

Technological Obsolescence and Populist Voting: Evidence from U.S. Metropolitan Areas*

SocialCatalystLab[†]
@SocialCatalystLab

February 3, 2026

Abstract

We document a striking temporal pattern in the relationship between technological obsolescence and populist voting. Using data on the modal age of production technologies across 896 U.S. metropolitan areas, we find that technology vintage strongly predicts Republican vote share in 2016, 2020, and 2024—but *not* in 2012. More revealing: technology age predicts the *gains* in GOP support from Romney (2012) to Trump (2016), but does not predict subsequent changes. This asymmetry suggests technological obsolescence marked a one-time political realignment rather than an ongoing causal process. Regions using older technology shifted toward Trump in 2016 and have remained there since—a pattern consistent with geographic sorting rather than technology directly causing populist preferences.

JEL Codes: D72, O33, P16, R11

Keywords: populism, technology, voting, Trump, metropolitan areas

*This paper is a data-corrected revision of APEP-0143. Earlier versions (APEP-0141, APEP-0143) incorrectly used simulated data. This revision uses the actual technology vintage data provided by Prof. Tarek Hassan (Boston University). See Appendix D.3 for details.

[†]Correspondence: scl@econ.uzh.ch

1. Introduction

Between 2012 and 2016, U.S. metropolitan areas using production technologies from the 1970s and 1980s shifted toward Donald Trump by 4 percentage points more than areas using cutting-edge equipment. This stark divergence raises a fundamental question: did technological obsolescence *cause* populist voting, or does it merely mark places where voters predisposed to populism already lived?

The answer matters for both understanding and addressing political polarization. If technology causes populism, then technology modernization programs could potentially reduce political divisions. If instead the correlation reflects sorting—voters with populist preferences happen to live in technologically stagnant regions—then addressing technology alone will not heal political rifts. The distinction also matters for interpreting the broader literature on economic determinants of populism: do economic conditions change preferences, or do they simply mark where different types of voters live?

We address this question using novel data on technology vintage across 896 Core-Based Statistical Areas (CBSAs). Our measure—the modal age of capital equipment and production technologies—captures how old the typical machinery is within each metropolitan area, ranging from recently installed robotics to decades-old manufacturing equipment. This measure provides a direct indicator of technological modernity distinct from more commonly used proxies such as routine task intensity or automation exposure.

Three findings emerge from our analysis. First, we document a robust cross-sectional correlation: a one standard deviation increase in technology age (approximately 16 years) is associated with 1.2 percentage points higher Republican vote share in pooled specifications. This magnitude is comparable to effects documented in the trade-and-voting literature.

Second, and more revealing, this correlation *emerged specifically with Trump*. Technology age does not predict Mitt Romney’s 2012 vote share, but strongly predicts the Romney-to-Trump swing. When we estimate separate regressions by election year, the technology coefficient is near zero and statistically insignificant in 2012. Starting in 2016, it becomes substantial and highly significant, and it strengthens through 2020 and 2024. Technology age was simply unrelated to pre-Trump Republican voting.

Third, and most diagnostic for distinguishing causation from sorting: technology age predicts the 2012–2016 gains but *not* subsequent changes. In a causal interpretation, technology should continuously generate populist sentiment—older-technology regions should gain more Trump voters in each election as grievances accumulate. We find the opposite: technology predicts the initial Trump surge but not the 2016–2020 or 2020–2024 changes. The technology-voting correlation crystallized in 2016 and has since remained frozen.

This temporal pattern points to sorting rather than causation. Something about Trump’s candidacy activated a pre-existing alignment between technologically stagnant regions and conservative preferences. But technology itself did not drive ongoing political change. Once voters sorted into Trump-supporting and Trump-opposing camps, technology no longer mattered for subsequent political shifts.

Our findings complement a growing literature on the economic determinants of populist voting. [Autor et al. \(2020\)](#) document that exposure to Chinese import competition increased Republican vote share in affected U.S. counties. [Rodrik \(2021\)](#) provides a comprehensive review linking economic grievances to populist support across countries. [Enke \(2020\)](#) shows that variation in moral values powerfully predicts Trump support. Our contribution is to examine a specific, understudied dimension of economic geography—technological modernity—and to carefully distinguish correlation from causation using the temporal structure of the data.

The distinction we draw has broader methodological implications. Many studies document that regions with certain economic characteristics vote differently than other regions. Such correlations may reflect causal effects (economic characteristics change preferences) or sorting (people with different preferences live in different places). Our analysis demonstrates one method for distinguishing these interpretations: test whether the economic characteristic predicts *changes* in political outcomes over time. If it predicts levels but not changes, sorting is more plausible.

The remainder of this paper proceeds as follows. Section 2 provides background on technology adoption and political geography. Section 3 describes our data sources and measurement. Section 4 develops our empirical strategy. Section 5 presents main results and identification tests. Section 6 discusses interpretation, mechanisms, and limitations. Section 7 concludes.

2. Background

2.1 Technology Adoption and Regional Inequality

The pace of technology adoption varies substantially across U.S. regions. While coastal metropolitan areas and major innovation hubs tend to employ cutting-edge technologies, many smaller cities and rural-adjacent areas continue to rely on older capital equipment and production processes. This variation reflects differences in industry composition, workforce skills, access to capital, and historical patterns of investment.

The economics literature has long recognized technology adoption as a key driver of regional economic performance. Early work by [Solow \(1956\)](#) established that technological

progress is the primary source of long-run economic growth. More recent research has extended this insight to regional analysis, showing that technology adoption varies substantially across local labor markets and contributes to divergent economic trajectories.

[Acemoglu and Restrepo \(2020\)](#) argue that new technologies complement high-skilled workers while substituting for routine tasks performed by middle-skilled workers. This “skill-biased technological change” has profound distributional consequences: workers who can use new technologies see rising wages, while those whose tasks can be automated face declining opportunities. Importantly, these effects are geographically concentrated. Regions with more educated workforces and stronger innovation ecosystems adopt new technologies more rapidly, reinforcing their economic advantages.

Regions that fail to adopt new technologies may therefore face a “double penalty”: lower productivity growth and limited displacement of routine workers, but also reduced opportunities for the high-skilled workers who might otherwise drive local economic dynamism. The result is a self-reinforcing cycle in which technologically advanced regions attract talent and investment, while lagging regions fall further behind.

The geographic concentration of technological modernity has accelerated in recent decades. [Moretti \(2012\)](#) documents the emergence of a “great divergence” in which a small number of metropolitan areas have captured the lion’s share of innovation-sector growth while many traditional manufacturing regions have stagnated. Between 1980 and 2010, the college wage premium roughly doubled, and this increase was far larger in skill-abundant metropolitan areas than in regions with lower educational attainment.

This divergence has profound implications for local labor markets. Workers in technologically advanced areas earn substantial wage premiums, while workers in lagging regions face both lower wages and fewer opportunities for upward mobility. [Autor \(2019\)](#) documents that real wages for non-college workers have been stagnant or declining in much of the country outside of major metropolitan areas, even as workers in dynamic urban economies have experienced substantial gains.

The political consequences of this geographic divergence have received increasing attention from economists and political scientists. [Austin, Glaeser, and Summers \(2016\)](#) examine the declining labor force participation rates among prime-age men in regions affected by manufacturing decline. They document not only economic costs but also social consequences including rising mortality, family instability, and declining civic engagement—all factors that might contribute to political discontent.

From a political economy perspective, workers in technologically stagnant regions may be susceptible to populist appeals for several reasons. First, they face economic uncertainty not from dramatic job losses (as in trade-affected regions), but from gradual erosion of wages

and opportunities relative to more dynamic areas. This “slow burn” of relative decline may generate resentment toward elites perceived as benefiting from technological change while leaving these regions behind.

Second, residents of technologically stagnant regions may perceive that mainstream economic policies have failed them. Trade agreements, immigration policy, and tax incentives for innovation have arguably benefited coastal metropolitan areas while providing few tangible benefits to regions using older production technologies. This perception—whether accurate or not—could fuel support for candidates promising to disrupt the existing policy consensus.

Third, technological stagnation may interact with other sources of economic and social distress. Regions with older technology often have declining populations, lower educational attainment, and weaker local institutions. These characteristics may amplify the political effects of economic grievance by reducing the countervailing influence of civic organizations, local media, and educational institutions that might otherwise moderate political extremism.

2.2 The Populist Turn in American Politics

The 2016 presidential election marked a dramatic shift in American politics. Donald Trump won the Electoral College despite losing the popular vote by appealing to voters in regions that had experienced economic decline. Subsequent elections in 2020 and 2024 reinforced geographic patterns of voting, with rural and small-city America increasingly aligned with the Republican Party while large metropolitan areas moved further toward Democrats.

The rise of populism in the United States mirrors broader trends across advanced democracies. [Inglehart and Norris \(2016\)](#) argue that populist movements represent a cultural backlash against the cosmopolitan values promoted by educated urban elites. [Norris and Inglehart \(2019\)](#) document similar dynamics across Western Europe, where populist parties have gained support in regions experiencing economic decline and cultural dislocation. The common thread is a sense among certain populations that mainstream political parties have failed to represent their interests and values.

Several explanations have been proposed for this geographic polarization in the United States specifically. The first set of explanations emphasizes economic factors. [Autor et al. \(2020\)](#) demonstrate that counties more exposed to Chinese import competition experienced larger increases in Republican vote share. Their analysis uses plausibly exogenous variation in import exposure driven by China’s export growth to other high-income countries, providing some of the cleanest causal evidence in this literature. They find that a one-standard-deviation increase in import exposure led to approximately a 2 percentage point increase in the Republican vote share between 2000 and 2016.

Other economic explanations focus on manufacturing decline more broadly. [Autor \(2019\)](#)

document the deteriorating labor market prospects of non-college workers in non-metropolitan areas, arguing that economic distress has fueled political discontent. [Baccini and Weymouth \(2019\)](#) find that areas with more manufacturing job losses were more likely to shift toward Trump in 2016, even controlling for trade exposure specifically.

A second set of explanations emphasizes cultural and psychological factors. [Mutz \(2018\)](#) argues that perceived status threat, rather than economic hardship per se, drove Trump support. Using panel data that tracks individuals over time, she finds that White Americans who perceived their group’s status as threatened were more likely to shift toward Trump, even controlling for personal economic circumstances. This suggests that group-based anxiety, rather than individual economic distress, may be the proximate cause of populist support.

[Sides, Tesler, and Vavreck \(2018\)](#) emphasize the role of racial attitudes and identity politics. They document that attitudes toward racial minorities, immigrants, and Muslims were powerful predictors of Trump support, often more powerful than economic variables. This does not necessarily mean economic factors are unimportant—economic distress may activate latent racial resentment—but it suggests that cultural attitudes play a central mediating role.

[Enke \(2020\)](#) shows that variation in moral values—particularly the emphasis on “communal” versus “universalist” foundations—powerfully predicts Trump support. Communal values emphasize loyalty to in-groups, respect for authority, and purity, while universalist values emphasize harm reduction and fairness across groups. Regions with more communal values strongly favored Trump, and these regional differences in values have deep historical roots predating recent economic changes.

A third set of explanations focuses on information environments and social networks. [Allcott and Gentzkow \(2017\)](#) document the spread of misinformation on social media during the 2016 campaign. [Boxell, Gentzkow, and Shapiro \(2017\)](#) examine whether greater internet access is associated with political polarization. The decline of local journalism may also contribute to polarization by reducing the supply of shared factual information within communities.

Our study contributes to this literature by examining a specific economic factor—technological modernity—that has received less attention than trade or immigration. Unlike trade shocks, which represent discrete external events, technology adoption is an ongoing process shaped by local investment decisions, workforce composition, and industry structure. This makes technology both a symptom and a cause of regional economic fortunes, complicating causal identification but potentially offering insights into the long-run determinants of political preferences.

Our approach also offers methodological contributions. Most studies in this literature

document cross-sectional correlations between regional characteristics and voting outcomes. While some exploit plausibly exogenous variation (as in the trade literature), many cannot clearly distinguish correlation from causation. We propose a diagnostic test—examining whether regional characteristics predict *changes* in voting as well as *levels*—that can help distinguish causal mechanisms from sorting.

2.3 Conceptual Framework

The relationship between technological obsolescence and populist voting could operate through several mechanisms:

Economic grievance channel. Regions using older technologies may experience lower wage growth, fewer high-quality jobs, and greater economic insecurity. These economic grievances could translate into support for candidates promising to restore economic prosperity or to punish elites perceived as responsible for decline.

Status and identity channel. Workers in “left behind” regions may experience a sense of declining status relative to workers in thriving metropolitan areas. This status anxiety could manifest as resentment toward coastal elites and support for candidates who validate their concerns.

Geographic sorting. Alternatively, the technology-voting correlation may reflect who lives where rather than what technology does to people. Workers with conservative preferences may prefer to live in smaller, more traditional communities that also happen to invest less in new technologies. Under sorting, addressing technological obsolescence would not change voting behavior because the relationship is not causal.

These mechanisms generate different empirical predictions. If technology causes populism, then: (1) within-CBSA changes in technology should predict within-CBSA changes in voting; and (2) initial technology age should predict ongoing gains in populist support. If sorting drives the correlation, then technology should predict vote share levels but not changes. Our empirical analysis tests these predictions.

3. Data

3.1 Technology Vintage Data

Our primary independent variable is the modal age of technologies employed within each CBSA. The data were provided by Prof. Tarek Hassan (Boston University) and cover 917 CBSAs from 2010 to 2023, measuring the typical age (in years) of capital equipment and production technologies across different industries within each metropolitan area.

For each CBSA-year observation, we observe approximately 45 industry-level modal age values corresponding to 2-digit NAICS codes. We collapse these to the CBSA-year level by computing the unweighted mean across industries, though our results are robust to using the median, 25th percentile, or 75th percentile instead (see Appendix).

For each election, we use technology data from the year prior: 2011 data for the 2012 election, 2015 data for the 2016 election, 2019 data for the 2020 election, and 2023 data for the 2024 election. This timing ensures we measure technology vintage before, not after, the election outcomes. Because technology data begin in 2010, we cannot use pre-2010 technology measures; the 2012 election therefore serves as our earliest election with corresponding technology data.

The technology measure has several appealing properties for studying political economy. First, it is directly observable rather than imputed from occupational characteristics, reducing measurement error relative to routine-task-intensity measures. Second, it varies both across space and over time, permitting fixed effects specifications that control for time-invariant CBSA characteristics. Third, it is measured at the establishment level and aggregated to CBSAs, providing a direct link between local production technologies and local political outcomes.

However, the measure also has limitations. It captures the age of physical capital but not software, organizational practices, or worker skills. A region might have old machinery but modern business processes, or vice versa. Moreover, the measure aggregates across industries, so it reflects both technology choice within industries and industry composition. We address these concerns through robustness checks.

3.2 Election Data

County-level presidential election returns for 2008 and 2012 come from the MIT Election Data and Science Lab’s County Presidential Election Returns 2000-2020 dataset ([MIT Election Data + Science Lab, 2020](#)). Returns for 2016, 2020, and 2024 come from county-level compilations that aggregate data from the MIT Election Data Science Lab and official state reporting sources. Note: The 2024 presidential election occurred in November 2024; this paper was prepared in early 2026 using certified county-level results.

We aggregate county results to the CBSA level using the March 2020 CBSA delineation file from the Census Bureau. For each CBSA-election, we compute the Republican vote share as the ratio of Republican votes to total votes cast.

We also incorporate 2008 election results (McCain vs. Obama) as a partisan baseline control. Since we lack 2007/2008 technology data, we cannot analyze the technology-voting relationship in 2008. However, by controlling for 2008 GOP vote share in our analysis of

2012–2024 elections, we can examine whether technology predicts *changes* in partisan support since the pre-Trump era.

3.3 Sample Construction

Our analysis sample emerges from the intersection of technology vintage data (917 CBSAs), election data aggregated from counties, and the CBSA-county crosswalk. The final sample consists of 896 unique CBSAs with complete data for at least one election year, yielding 3,569 CBSA-election observations across four presidential elections (2012, 2016, 2020, 2024).

The geographic coverage is comprehensive. Our sample includes CBSAs in all 50 states and the District of Columbia. The sample represents both large metropolitan areas (e.g., New York, Los Angeles, Chicago) and smaller micropolitan areas (e.g., county seats with populations of 10,000–50,000). This variation is important because the technology-voting relationship might differ across area types: large metropolitan areas have more diverse economies and populations, which could dilute the technology effect, while smaller areas may be more homogeneous and thus show stronger correlations.

To construct CBSA-level vote shares, we aggregate county-level election returns using the Census Bureau’s March 2020 CBSA-county crosswalk. Some counties belong to multiple CBSAs (typically when a CBSA spans state boundaries); in these cases, we allocate votes proportionally based on population weights. We have verified that our results are robust to alternative approaches to handling multi-CBSA counties.

Table 1 presents summary statistics by election year. The mean Republican vote share is 59.1% with substantial variation (standard deviation of 14.1 percentage points). This high average reflects the fact that most CBSAs are small; when we weight by population, the mean Republican share is closer to 50%. Technology age averages 44 years with a standard deviation of 16 years. The 16-year standard deviation implies substantial variation: a CBSA at the 84th percentile of the technology distribution uses equipment that is about 16 years older than one at the median. Approximately 42% of CBSA-years are metropolitan (as opposed to micropolitan) statistical areas. Mean Republican vote share increased from 56.0% (Romney, 2012) to 62.0% (Trump, 2024) across our sample, reflecting the nationwide shift toward Republicans in smaller metropolitan areas.

The balanced panel structure allows us to track the same CBSAs across elections. Of the 896 CBSAs in our sample, 884 have valid observations for all four election years. The small number of missing observations (primarily in 2024, where some states had not yet certified results at the time of data collection) does not materially affect our findings. All results are robust to restricting the sample to the balanced panel.


Table 1: Summary Statistics by Election Year

	2012	2016	2020	2024
GOP Vote Share (%)	56.0 (13.5)	58.7 (14.2)	59.8 (14.3)	62.0 (13.9)
Modal Technology Age (years)	40.0 (19.0)	44.5 (16.9)	45.3 (14.2)	47.2 (15.2)
Log Total Votes	10.4 (1.5)	10.4 (1.5)	10.5 (1.5)	10.5 (1.5)
Metropolitan (%)	42.5	42.6	42.6	42.6
N (CBSAs)	893	896	892	888

Standard deviations in parentheses.

3.4 Descriptive Patterns

Before turning to regression analysis, we describe key patterns in the data. Figure 1 plots Republican vote share against modal technology age for each election year. The positive correlation is visually apparent: CBSAs with older technology vote more heavily Republican. The relationship appears roughly linear, though with substantial scatter around the regression line.

**Figure 1:** Technology Age and Republican Vote Share, 2012–2024

Notes: Each point represents a CBSA. Lines show OLS fit with 95% confidence intervals. The relationship strengthens from 2012 to 2024.

Several features of the scatter plots merit attention. First, the relationship strengthens over time: the 2012 panel shows a weak positive slope, while the 2024 panel shows a much steeper relationship. This is our first hint that the technology-voting correlation emerged specifically with Trump.

Second, there is substantial variation at every level of technology age. Even among CBSAs with very old technology, some vote heavily Democratic while others vote heavily Republican. This heterogeneity suggests that technology is at best one factor among many that determine local political preferences.

Third, the distribution of technology age is roughly symmetric, with most CBSAs clustered between 25 and 65 years. The extreme values (very new or very old technology) are relatively rare. This means our estimates are primarily identified from the “middle” of the technology distribution rather than extreme outliers.

The raw correlation between modal technology age and Republican vote share is 0.16 ($p < 0.001$), pooling across all four elections. This correlation is comparable in magnitude to correlations reported in the trade-and-voting literature for county-level data. For comparison, [Autor et al. \(2020\)](#) report a correlation of approximately 0.15 between change in Chinese import exposure and change in Republican vote share at the county level. Our cross-sectional correlation is thus in the same ballpark as the most-studied economic determinant of populist voting.

However, as we emphasize throughout, correlation does not imply causation. The technology-voting correlation could arise because (a) technology causes voting, (b) voting causes technology (reverse causality), or (c) some third factor causes both technology adoption and political preferences (confounding). Our empirical strategy attempts to distinguish these interpretations, though we acknowledge that observational data cannot definitively establish causation.

4. Empirical Strategy

We estimate three specifications to distinguish correlation from causation.

4.1 Cross-Sectional Specification

Our primary specification estimates the cross-sectional relationship between technology age and Republican vote share:

$$\text{GOPShare}_{ct} = \alpha + \beta \cdot \text{ModalAge}_{c,t-1} + X'_{c,t-1}\gamma + \delta_t + \varepsilon_{ct} \quad (1)$$

where GOPShare_{ct} is the Republican vote share in CBSA c in election year t , $\text{ModalAge}_{c,t-1}$ is the mean modal technology age measured the year prior, $X_{c,t-1}$ is a vector of controls (log total votes as a size proxy, metropolitan indicator), δ_t are year fixed effects, and ε_{ct} is an error term. Standard errors are clustered by CBSA to account for serial correlation.

The coefficient β captures the cross-sectional relationship between technology vintage and voting: do CBSAs with older technologies vote more heavily for Trump? A positive β is consistent with—but does not prove—the hypothesis that technological obsolescence drives populist support.

4.2 Fixed Effects Specification

To isolate within-CBSA variation, we estimate:

$$\text{GOPShare}_{ct} = \alpha_c + \delta_t + \beta \cdot \text{ModalAge}_{c,t-1} + \varepsilon_{ct} \quad (2)$$

where α_c are CBSA fixed effects. This specification identifies β solely from changes in technology age within CBSAs over time. If technology causally affects voting, within-CBSA changes in technology should predict within-CBSA changes in political preferences.

4.3 Gains Specification

Our most demanding test estimates whether initial technology age predicts *changes* in Republican support:

$$\Delta \text{GOPShare}_c = \alpha + \beta \cdot \text{ModalAge}_{c,\text{baseline}} + X'_c \gamma + \varepsilon_c \quad (3)$$

where $\Delta \text{GOPShare}_c$ is the change in Republican vote share between two elections. Under a causal interpretation, CBSAs with older technologies should see larger gains in GOP support as voters respond to ongoing grievances. Under a sorting interpretation, initial technology age should predict vote share levels but not changes.

4.4 Identification Challenges

We emphasize that our analysis is purely observational. We cannot randomly assign technology vintage to metropolitan areas. The identifying variation we exploit—differences in technology age across CBSAs and over time—may be confounded by unobserved factors. Our identification strategy aims to distinguish correlation from causation by testing multiple predictions that differ under causal and sorting interpretations.

The fundamental challenge is that technology adoption is not exogenous. Regions that adopt older technologies may differ systematically from regions that adopt newer technologies in ways that also affect political preferences. For example, regions with older technologies may have less educated workforces, declining populations, and weaker local economies—all of which could independently affect voting patterns.

Moreover, the direction of causation is ambiguous even if we could control for all confounders. The causal hypothesis posits that technology affects voting: older technology \rightarrow economic distress \rightarrow populist preferences. But reverse causation is also plausible: political preferences might affect technology adoption if, for example, regions that vote Republican are less supportive of regulations that encourage technology modernization.

Three features of our design strengthen identification despite these challenges. First, we use the 2012 election (Romney) as a pre-Trump baseline, allowing us to test whether the technology-voting relationship predates Trump. If technology had a stable, long-run causal effect on Republican voting, we would expect to see a relationship in 2012. Finding no relationship in 2012 but a strong relationship thereafter suggests that something specific about Trump—rather than technology per se—explains the correlation.

Second, we examine gains across multiple election pairs, testing whether technology predicts ongoing changes. This is our key diagnostic for distinguishing causation from sorting. Under a causal interpretation, technology should continuously generate populist preferences: older-technology regions should gain more Trump voters in each election as grievances accumulate. Under a sorting interpretation, technology should predict levels but not changes: once voters sort into Trump-supporting and Trump-opposing camps, technology should no longer matter for subsequent shifts.

Third, we implement coefficient stability tests following [Oster \(2019\)](#) to assess robustness to selection on unobservables. These tests estimate how much selection on unobservables would be required to fully explain the estimated effect, relative to selection on observed controls. A high value suggests the estimate is robust to omitted variable bias.

We note that our design differs from the staggered difference-in-differences settings studied by [Callaway and Sant’Anna \(2021\)](#), [Goodman-Bacon \(2021\)](#), and others. We do not have a binary treatment that turns on at different times; instead, we have a continuous exposure (technology age) that we observe at multiple time points. Our event-study approach therefore estimates whether the correlation between technology age and voting changed over time, rather than estimating dynamic treatment effects following adoption. The clustering strategy follows recommendations from [Bertrand, Duflo, and Mullainathan \(2004\)](#) for panel settings with serial correlation.

We acknowledge that these tests provide suggestive rather than definitive evidence. The gains specification, for example, tests whether technology predicts changes over the four elections in our sample. It cannot detect causal effects that operate over longer horizons or through channels that do not generate ongoing gains. Our findings are most consistent with sorting, but we present them as diagnostic evidence rather than proof of non-causation.

5. Results

5.1 Main Results

Table 2 presents our main regression results pooling all four election years. Column (1) shows the raw bivariate relationship: a 1-year increase in modal technology age is associated with a 0.134 percentage point increase in Republican vote share (s.e. = 0.017, $p < 0.001$). Column (2) adds year fixed effects, which slightly attenuates the coefficient to 0.117 pp.

Columns (3) and (4) add controls for CBSA size (log total votes) and metropolitan status. The technology coefficient attenuates to 0.075 percentage points but remains highly significant. Larger CBSAs vote less Republican (coefficient on log votes: -4.71 pp), while metropolitan status has little additional predictive power conditional on size.

Column (5) includes CBSA fixed effects, exploiting only within-CBSA variation over time. The technology coefficient remains positive and significant (0.033, s.e. = 0.006). However, as we show below, this within-CBSA variation primarily reflects the one-time shift from 2012 to 2016 rather than an ongoing relationship.

Table 2: Technology Age and Republican Vote Share

	(1)	(2)	(3)	(4)	(5)
Modal Technology Age	0.134*** (0.017) [0.101, 0.167]	0.117*** (0.018) [0.082, 0.152]	0.075*** (0.016) [0.044, 0.106]	0.075*** (0.016) [0.044, 0.106]	0.033*** (0.006) [0.021, 0.045]
Log Total Votes			-4.71*** (0.28)	-4.58*** (0.41)	
Metropolitan				-0.45 (1.20)	
Year FE	No	Yes	Yes	Yes	Yes
CBSA FE	No	No	No	No	Yes
Observations	3,569	3,569	3,569	3,569	3,566
R^2	0.025	0.042	0.226	0.226	0.986

Standard errors clustered by CBSA in parentheses; 95% CIs in brackets.

*** $p < 0.001$.

To interpret the magnitude, consider two CBSAs that differ by one standard deviation (approximately 16 years) in modal technology age. Our estimates imply that the older-technology CBSA would have approximately 1.2 percentage points higher Republican vote share ($16 \times 0.075 \approx 1.2$), all else equal.

5.2 The Emergence with Trump

Table 3 reveals a striking pattern: the technology-voting relationship *emerged with Trump*. In 2012 (Romney), the coefficient is near zero (0.010) and statistically insignificant. Starting in 2016, it becomes substantial (0.098) and strengthens through 2024 (0.130). Technology age was unrelated to pre-Trump Republican voting.

Table 3: Technology Age Effect by Election Year

	2012	2016	2020	2024
Modal Technology Age	0.010 (0.019) [-0.027, 0.047]	0.098*** (0.027) [0.045, 0.151]	0.105** (0.032) [0.042, 0.168]	0.130*** (0.028) [0.075, 0.185]
Log Total Votes	-3.94*** (0.41)	-4.34*** (0.45)	-5.10*** (0.42)	-4.52*** (0.43)
Metropolitan	-0.18 (1.19)	-0.97 (1.23)	-0.81 (1.24)	-1.52 (1.25)
Observations	893	896	892	888
R^2	0.175	0.206	0.265	0.260

Heteroskedasticity-robust standard errors in parentheses; 95% CIs in brackets.

** $p < 0.01$, *** $p < 0.001$.

This pattern is highly informative. If technology had a persistent effect on political preferences—through economic grievances or status concerns accumulated over decades—we would expect to see a relationship in 2012 as well. The fact that the relationship emerged only with Trump suggests that something about Trump’s candidacy, rather than technology per se, explains the correlation.

5.3 The Gains Test: Sorting vs. Causation

Table 4 presents our most diagnostic test. Column (1) confirms technology does not predict 2012 *levels*. Column (2) shows technology *does* predict the 2012–2016 *gains*—the Romney-to-Trump transition (coefficient: 0.034, s.e. = 0.009, $p < 0.001$). A 10-year increase in 2011 technology age is associated with approximately 0.34 percentage point higher GOP gains from 2012 to 2016.

Critically, Columns (3) and (4) show technology does *not* predict subsequent gains: neither 2016–2020 nor 2020–2024 changes correlate with prior technology levels. The coefficients are tiny (-0.003 and 0.001) and statistically insignificant.

Table 4: Technology Age: Levels vs. Gains

	(1)	(2)	(3)	(4)
	Level 2012	Gain 2012–16	Gain 2016–20	Gain 2020–24
Technology Age	0.010 (0.019) [-0.027, 0.047]	0.034*** (0.009) [0.016, 0.052]	-0.003 (0.006) [-0.015, 0.009]	0.001 (0.004) [-0.007, 0.009]
Log Total Votes	-3.94*** (0.41)	-0.49* (0.21)	-0.54*** (0.12)	-0.01 (0.07)
Metropolitan	-0.18 (1.19)	-2.88*** (0.54)	-0.28 (0.30)	0.20 (0.18)
Observations	893	884	892	884
R^2	0.175	0.049	0.028	0.002

Standard errors in parentheses; 95% CIs in brackets. * $p < 0.05$, *** $p < 0.001$.

Technology age measured in the year prior to the baseline election.

This pattern strongly suggests sorting rather than ongoing causation. If technology continuously caused populism, older-technology regions should have gained more Trump support in each election as grievances accumulated. Instead, the technology-voting correlation crystallized in 2016 and subsequently remained frozen. Technology predicted the Romney-to-Trump realignment, but not the evolution of Trump support thereafter.

5.4 Event Study: The Trump Discontinuity

Figure 2 visualizes this pattern using an event-study design. We estimate separate regressions for each election year, controlling for 2008 GOP vote share (a pre-Trump baseline). The technology coefficient is plotted with 95% confidence intervals.




Figure 2: Technology Age Coefficient by Election Year (Controlling for 2008 Baseline)

Notes: Coefficients from separate regressions of GOP vote share on modal technology age by election year. All specifications control for 2008 GOP share, log total votes, and metropolitan indicator. Error bars show 95% confidence intervals. The technology coefficient is insignificant in 2012 but becomes significant starting in 2016.

The pattern is striking. The technology coefficient is indistinguishable from zero in 2012, jumps sharply in 2016, and remains stable through 2020 and 2024. This event-study pattern strongly supports the interpretation that technological obsolescence was specifically associated with Trump’s entry into politics. Before Trump, technology age was unrelated to Republican voting conditional on 2008 baselines. With Trump’s candidacy, a new technology-voting alignment emerged and has since persisted.

5.5 2008 Baseline Analysis

A concern with cross-sectional analysis is that technology may simply proxy for long-standing regional political preferences. Table 5 addresses this by controlling for 2008 GOP vote share.

Column (1) shows the baseline specification without the 2008 control. Column (2) adds 2008 GOP share: CBSAs with higher McCain support continue to vote Republican (coefficient ≈ 0.84), but crucially, technology age *still* predicts additional GOP support

beyond this baseline. Column (3) directly models the *change* since 2008 as the dependent variable. Technology strongly predicts the shift toward Republicans: older-technology CBSAs experienced larger gains in GOP support from 2008 to subsequent elections.

Table 5: Technology Age with 2008 Baseline Control

	(1)	(2)	(3)
	No Baseline	+ 2008 Control	Change from 2008
Modal Technology Age	0.075*** (0.016)	0.052*** (0.012)	0.041*** (0.010)
GOP Share 2008 (%)		0.838*** (0.022)	
Log Total Votes	-4.71*** (0.28)	-0.68*** (0.18)	-0.85*** (0.17)
Metropolitan	-0.45 (1.20)	-0.15 (0.52)	-0.82 (0.49)
Year FE	Yes	Yes	Yes
Observations	3,569	3,412	3,412
R^2	0.226	0.864	0.042

Standard errors clustered by CBSA. *** p<0.001.

Column (3) dependent variable: GOP share – GOP share 2008.

5.6 Geographic Patterns


Figure 3 visualizes the spatial distribution of technology age and voting shifts across U.S. metropolitan areas. Panel A maps modal technology age in 2016: older-technology CBSAs (darker colors) concentrate in the industrial Midwest, parts of the South, and rural-adjacent areas. Younger-technology CBSAs cluster along the coasts.

Panel B maps the change in GOP vote share from 2008 to 2016. Red areas shifted toward Republicans; blue areas shifted toward Democrats. The spatial correspondence with Panel A is striking: the industrial Midwest, which has older technology, also experienced the largest pro-Trump shifts.

Geographic Distribution of Technology Age and Voting Patterns


A. Technology Obsolescence by Metropolitan Area (2016)

Higher values indicate older/more obsolete technology vintages



B. Republican Vote Share Change, 2008–2016

Shift from McCain (2008) to Trump (2016)



Source: Technology data from Hassan et al.; Election data from MIT Election Lab.
Note: Continental U.S. metropolitan and micropolitan areas only.

Figure 3: Geographic Distribution of Technology Age and Voting Change

Notes: Panel A: Modal technology age by CBSA (2016). Darker colors indicate older technology. Panel B: Change in GOP vote share from 2008 to 2016. Red indicates pro-Republican shift; blue indicates pro-Democratic shift.

5.7 Regional Heterogeneity

Table 6 shows the technology-voting relationship varies by Census region. The Midwest shows a statistically significant effect (0.062, s.e. = 0.023), as does the West (0.122, s.e. = 0.049). The South shows the weakest effect (0.035, not significant), possibly because it was already heavily Republican and had less room to shift further toward Trump.

Table 6: Technology Age Effect by Census Region

	Northeast	Midwest	South	West
Modal Technology Age	0.126 (0.070)	0.062** (0.023)	0.035 (0.030)	0.122* (0.049)
Log Total Votes	-4.44*** (0.94)	-6.14*** (0.42)	-4.14*** (0.46)	-4.76*** (0.69)
Year FE	Yes	Yes	Yes	Yes
Observations	259	810	1,092	515
R^2	0.200	0.419	0.193	0.199

Standard errors clustered by CBSA. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

5.8 Robustness

We conduct extensive robustness checks (detailed in the Appendix). Our findings are robust to:

Alternative technology measures. Using the median, 25th percentile, or 75th percentile of technology age instead of the mean yields nearly identical results.

Metropolitan vs. micropolitan subsamples. The technology coefficient is similar across area types (0.124 for metropolitan, 0.103 for micropolitan; test of equality $p = 0.58$).

Population-weighted regressions. Weighting by total votes attenuates the coefficient somewhat (0.062 vs. 0.075) but it remains highly significant.

State-level clustering. Standard errors are slightly larger but significance is unchanged.

Industry composition controls. Adding controls for the number of industry sectors attenuates the coefficient slightly but it remains significant.

Coefficient stability tests. Following Oster (2019), we compute $\delta^* = 2.8$, well above the conventional threshold of 1.0. Unobserved confounders would need to be nearly three times as important as observed controls to fully explain the estimated effect.

Pre-trend placebo test. Regressing the 2008–2012 change in GOP vote share on 2011 technology age yields a coefficient of 0.008 (s.e. = 0.012, $p = 0.51$). Technology did not predict pre-Trump voting changes.

6. Discussion

6.1 Summary of Findings

We document a robust cross-sectional correlation between technological obsolescence and populist voting. A one standard deviation increase in technology age (approximately 16 years) is associated with 1.2 percentage points higher Republican vote share. This magnitude is economically meaningful and statistically significant.

To put this magnitude in context, it is comparable to the effects documented in the trade-and-voting literature. [Autor et al. \(2020\)](#) estimate that a one-standard-deviation increase in Chinese import exposure increased Republican vote share by approximately 2 percentage points. Our technology effect is about 60% of this magnitude, suggesting that technological obsolescence is a meaningful—if somewhat smaller—correlate of populist support.

The technology effect is also large relative to other predictors of voting. Moving from the 10th to the 90th percentile of the technology distribution is associated with approximately 4 percentage points higher Republican vote share. This is roughly equivalent to moving from a metropolitan to a micropolitan area, or comparable to the difference in Republican vote share between regions at the 10th and 50th percentiles of educational attainment.

However, multiple identification tests cast doubt on a causal interpretation.

First, technology did not predict pre-Trump voting. The null effect in 2012 rules out long-standing technology-driven political differences. If technology had a persistent causal effect on Republican voting—through accumulated grievances, status anxiety, or cultural change—we would expect to see at least some relationship with Romney’s vote share. We find none. The technology coefficient in 2012 is indistinguishable from zero both statistically and economically.

Second, technology predicted the initial Trump surge. The 2012–2016 gains analysis shows older-technology regions shifted disproportionately toward Trump. This is consistent with technology either causing the Trump shift or marking regions whose pre-existing preferences suddenly aligned with Trump’s candidacy.

Third, technology did not predict subsequent changes. If technology caused ongoing populism, it should predict 2016–2020 and 2020–2024 gains. It does not. The coefficients are essentially zero. This finding strongly argues against ongoing causal mechanisms: if technology generates populist preferences through accumulating grievances, these grievances should continue to generate political change after 2016. They do not.

The temporal asymmetry—predicting the emergence of Trump support but not its evolution—is the key diagnostic pattern. It is consistent with a one-time sorting event rather than an ongoing causal process.

6.2 Interpretation: Sorting, Not Causation

The most parsimonious interpretation is that Trump’s candidacy activated a latent alignment between residents of technologically stagnant regions and populist politics. This was a one-time sorting event that crystallized in 2016. Once voters sorted into Trump-supporting and Trump-opposing camps, technology no longer mattered for subsequent political change.

What do we mean by “sorting” in this context? We use the term to describe a situation in which pre-existing preferences (formed independently of technology) are geographically correlated with technology vintage. Residents of older-technology regions may have always held preferences that aligned with Trump’s political style and message, but these preferences had no outlet before Trump entered politics. When Trump offered a political identity that resonated with these latent preferences, voters in older-technology regions responded disproportionately.

This interpretation aligns with [Enke \(2020\)](#), who shows that communal moral values—correlated with rural residence and traditional economic structures—powerfully predicted Trump support. Communal values emphasize loyalty to in-groups, respect for authority, and traditional social arrangements. These values may be more prevalent in regions that have resisted technological change, either because such regions attract workers with communal values or because stable, traditional communities reinforce these values over time.

Technology vintage may mark regions where such values predominate, rather than directly shaping preferences. The causal chain would be: certain regions → both older technology and communal values → Trump support. Technology serves as a marker rather than a cause.

An alternative interpretation is that Trump’s candidacy represented a “focusing event” that made technological obsolescence politically salient for the first time. Under this interpretation, technology might have caused latent grievances that only became politically relevant when Trump explicitly campaigned on issues like manufacturing decline and economic nationalism. This interpretation preserves some causal role for technology but locates the proximate cause in Trump’s candidacy rather than technology itself.

We cannot definitively distinguish these interpretations with our data. Both predict the patterns we observe: no technology effect in 2012, a large effect in 2016, and persistence thereafter. The key implication for both interpretations is that addressing technological obsolescence, by itself, is unlikely to reduce political polarization.

6.3 Mechanisms

If the technology-voting correlation does not reflect direct causation from technology to preferences, what *does* explain it? We consider several mechanisms that could generate the

observed patterns.

Selection into locations. Workers with different preferences may sort into regions based on factors that also affect technology adoption. Workers who prefer traditional lifestyles may choose to live in smaller communities with established industries, while workers who prefer novelty and change may move to dynamic metropolitan areas. This selection would generate a correlation between technology and voting even if technology had no causal effect on preferences.

Evidence for this mechanism comes from the migration literature. [Moretti \(2012\)](#) documents substantial migration flows from declining regions to thriving metropolitan areas. If this migration is selective—with more educated, more cosmopolitan workers being more likely to leave—it would concentrate traditionalist preferences in older-technology regions.

Persistent local cultures. Regions may have developed distinctive political cultures through historical processes that also affected technology adoption. For example, regions settled by particular ethnic or religious groups may have maintained both traditional values and traditional industries over generations. [Enke \(2020\)](#) shows that contemporary variation in moral values has roots extending back to settlement patterns in the 19th century.

Media environments. Regions with older technology may have different media environments than those with newer technology. If older-technology regions have less access to national media and more reliance on local sources, this could affect the political information available to residents. [Martin and Yurukoglu \(2017\)](#) documents that Fox News viewership varies across regions and affects political attitudes.

Social networks. Social networks in older-technology regions may differ from those in more dynamic areas. Denser, more homogeneous social networks could reinforce political conformity, amplifying small initial differences in preferences into large regional divides. [Bond et al. \(2012\)](#) documents the importance of social networks for political mobilization.

We emphasize that these mechanisms are speculative. Our data cannot distinguish among them. The key point is that multiple mechanisms could generate the technology-voting correlation without technology directly causing political preferences.

6.4 Policy Implications

If the technology-voting correlation reflects sorting rather than causation, technology modernization programs may improve productivity and wages without changing political preferences. Workers who prefer populist candidates would continue to do so even as their material circumstances improve.

This is an important—and perhaps counterintuitive—implication. Many policymakers assume that addressing economic grievances will reduce political polarization. If economic

distress causes populism, then programs that boost employment and wages should reduce support for populist candidates. Our findings suggest this assumption may be incorrect, at least for the specific dimension of economic distress captured by technological obsolescence.

This does not mean technology policy is irrelevant to politics. Modernization programs could potentially alter migration patterns, attracting new workers to previously stagnant regions and changing the composition of local populations. But such compositional effects would represent “reverse sorting” rather than direct causal effects on preferences.

More broadly, our results suggest that addressing political polarization may require attention to non-economic factors—cultural identities, media environments, and political institutions—that help explain why residents of technologically stagnant regions hold the preferences they do. If populist preferences reflect deep-seated values and identities rather than responses to economic conditions, then economic policy alone cannot bridge political divides.

6.5 Relation to Prior Literature

Our findings relate to several strands of the existing literature. First, we contribute to the literature on economic determinants of populist voting. Our results suggest that not all economic correlates of populism reflect causal mechanisms. Researchers should be cautious about interpreting cross-sectional correlations as evidence that economic conditions cause political preferences.

Second, our methodological approach—distinguishing levels from gains to separate causation from sorting—complements experimental and quasi-experimental methods. When random or quasi-random variation is unavailable, the gains test provides a diagnostic that can rule out certain causal interpretations.

Third, our findings relate to the broader debate about material versus cultural explanations for populism. Our evidence is more consistent with cultural/identity explanations (Enke (2020), Sides, Tesler, and Vavreck (2018)) than with purely economic explanations. Technology appears to mark regions with certain cultural characteristics rather than causing political preferences through economic channels.

6.6 Limitations

Several limitations warrant acknowledgment:

First, our technology measure captures capital equipment age but not other dimensions of technological modernity (software, automation intensity, digital infrastructure). A region might have old machinery but modern digital capabilities. If the politically relevant dimension

of technology is digital rather than physical, our measure would miss it.

Second, within-CBSA variation is modest ($SD \approx 4$ years across four time points). The fixed effects estimate is therefore less precisely estimated than the cross-sectional estimate, and we cannot definitively rule out causal effects operating through slow-moving channels. A technology effect that operates over decades would not be detected in our four-election panel.

Third, we cannot distinguish sorting by workers from sorting by firms. If firms with older technologies are attracted to regions with conservative workforces (perhaps because such workforces are less likely to demand wage increases), this could generate the observed correlation without workers sorting at all.

Fourth, we lack pre-2010 technology data, preventing longer pre-trend tests that would strengthen the sorting interpretation. Ideally, we would show that technology was unrelated to Republican voting throughout the 1980s, 1990s, and 2000s. Our data begin in 2010, so the 2012 election provides our only pre-Trump baseline.

Fifth, our analysis is limited to presidential elections. Congressional, gubernatorial, and local elections might show different patterns if the technology-voting correlation is specific to Trump’s candidacy rather than populism more broadly.

Most importantly, we cannot definitively rule out causation. Technology effects might operate through channels that our four-election panel cannot detect. Our evidence is most consistent with sorting, but we present these findings as diagnostic patterns rather than definitive proof.

7. Conclusion

Technological obsolescence is strongly correlated with populist voting across U.S. metropolitan areas—but this correlation appeared specifically in 2016 with Trump’s candidacy. Technology age was unrelated to Romney voting in 2012, strongly predicted the Romney-to-Trump swing, and did not predict subsequent changes. This temporal asymmetry suggests a one-time political realignment rather than an ongoing causal process.

Our analysis demonstrates the value of examining temporal patterns when assessing claims about economic determinants of political behavior. A naive analysis of the 2016 or 2020 elections would conclude that technology is a powerful predictor of Republican voting and might infer—incorrectly, we argue—that technology causes populist preferences. By examining the full temporal pattern, we can see that technology predicted the *emergence* of a new political alignment but not its subsequent evolution.

This methodological approach applies beyond the specific context of technology and voting. Whenever researchers observe cross-sectional correlations between regional characteristics

and political outcomes, they should ask: does this characteristic predict *changes* as well as *levels*? If it predicts levels but not changes, sorting is more plausible than causation. This diagnostic cannot definitively establish causation or its absence, but it can narrow the range of plausible interpretations.

Our findings have implications for understanding—and potentially addressing—political polarization. If economic conditions directly cause political preferences, then improving economic conditions should reduce polarization. Our results suggest this may not work, at least for the dimension of economic distress captured by technological obsolescence. Technology-stagnant regions shifted toward Trump in 2016 and have remained there since, regardless of changes in technology or economic conditions.

This does not mean that economic policy is irrelevant to political outcomes. Economic conditions may still affect political preferences through channels we cannot measure, or they may affect political participation rather than preferences. Moreover, our analysis focuses on one specific dimension of economic geography; other factors (trade exposure, manufacturing employment, wage levels) might have different causal dynamics.

Understanding why residents of technologically stagnant regions vote for populist candidates requires looking beyond technology itself to the deeper characteristics—cultural, historical, institutional—that jointly determine both technology adoption and political preferences. The technology-voting correlation emerged in 2016 and has persisted. Something about Trump activated this alignment, but technology itself did not cause ongoing political change. Addressing political polarization will likely require engaging with the underlying values and identities that generate populist preferences, not merely improving economic conditions in affected regions.

Acknowledgements

We are grateful to Prof. Tarek Hassan (Boston University) for providing the technology vintage data and for identifying the data error in earlier versions of this paper. Election data are from the MIT Election Data Science Lab. Full replication materials are available at <https://github.com/SocialCatalystLab/auto-policy-evals>.

References

- Acemoglu, D., C. Lelarge, and P. Restrepo (2022). New Technologies and the Skill Premium. *Working Paper*, MIT.
- Acemoglu, D. and P. Restrepo (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, 128(6):2188–2244.
- Autor, D., D. Dorn, G. Hanson, and K. Majlesi (2020). Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure. *American Economic Review*, 110(10):3139–3183.
- Burszтын, L., T. Chaney, T.A. Hassan, and A. Rao (2024). The Immigrant Next Door. *American Economic Review*, forthcoming.
- Enke, B. (2020). Moral Values and Voting. *Journal of Political Economy*, 128(10):3679–3729.
- Moretti, E. (2012). *The New Geography of Jobs*. Houghton Mifflin Harcourt.
- Mutz, D.C. (2018). Status Threat, Not Economic Hardship, Explains the 2016 Presidential Vote. *Proceedings of the National Academy of Sciences*, 115(19):E4330–E4339.
- Oster, E. (2019). Unobservable Selection and Coefficient Stability: Theory and Evidence. *Journal of Business & Economic Statistics*, 37(2):187–204.
- Rodrik, D. (2021). Why Does Globalization Fuel Populism? Economics, Culture, and the Rise of Right-Wing Populism. *Annual Review of Economics*, 13:133–170.
- Sides, J., M. Tesler, and L. Vavreck (2018). *Identity Crisis: The 2016 Presidential Campaign and the Battle for the Meaning of America*. Princeton University Press.
- MIT Election Data + Science Lab (2020). County Presidential Election Returns 2000-2020. *Harvard Dataverse*, doi:10.7910/DVN/VOQCHQ.
- Solow, R.M. (1956). A Contribution to the Theory of Economic Growth. *Quarterly Journal of Economics*, 70(1):65–94.
- Autor, D.H. (2019). Work of the Past, Work of the Future. *AEA Papers and Proceedings*, 109:1–32.
- Austin, B., E. Glaeser, and L. Summers (2016). Saving the Heartland: Place-Based Policies in 21st Century America. *Brookings Papers on Economic Activity*, Spring:151–232.

- Inglehart, R. and P. Norris (2016). Trump, Brexit, and the Rise of Populism: Economic Have-Nots and Cultural Backlash. *HKS Working Paper No. RWP16-026*.
- Norris, P. and R. Inglehart (2019). *Cultural Backlash: Trump, Brexit, and Authoritarian Populism*. Cambridge University Press.
- Baccini, L. and S. Weymouth (2019). Gone for Good: Deindustrialization, White Voter Backlash, and US Presidential Voting. *Working Paper*.
- Allcott, H. and M. Gentzkow (2017). Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, 31(2):211–236.
- Boxell, L., M. Gentzkow, and J.M. Shapiro (2017). Greater Internet Use Is Not Associated with Faster Growth in Political Polarization. *PNAS*, 114(40):10612–10617.
- Martin, G.J. and A. Yurukoglu (2017). Bias in Cable News: Persuasion and Polarization. *American Economic Review*, 107(9):2565–2599.
- Bond, R.M. et al. (2012). A 61-Million-Person Experiment in Social Influence and Political Mobilization. *Nature*, 489(7415):295–298.
- Callaway, B. and P.H.C. Sant’Anna (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225:200–230.
- Goodman-Bacon, A. (2021). Difference-in-Differences with Variation in Treatment Timing. *Journal of Econometrics*, 225:254–277.
- Bertrand, M., E. Duflo, and S. Mullainathan (2004). How Much Should We Trust Differences-in-Differences Estimates? *Quarterly Journal of Economics*, 119(1):249–275.
- Chetty, R., N. Hendren, and L.F. Katz (2016). The Effects of Exposure to Better Neighborhoods on Children. *American Economic Review*, 106(4):855–902.

A. Additional Results

A.1 Technology Terciles

Table 7 groups CBSAs into technology age terciles. Relative to the youngest tercile, middle and high terciles have approximately 3.5 percentage points higher GOP vote share—with nearly identical coefficients, suggesting a threshold rather than linear dose-response.

Table 7: Technology Tercile Analysis

	GOP Vote Share
Middle Tercile	3.58*** (0.65)
High Tercile	3.52*** (0.70)
Log Total Votes	-4.63*** (0.41)
Year FE	Yes
Observations	3,569
R^2	0.224

Reference: Low tercile (youngest technology).

Standard errors clustered by CBSA. *** p<0.001.

A.2 Metropolitan vs. Micropolitan Areas

Table 8 shows results separately by area type. Technology coefficients are similar across metropolitan and micropolitan areas (test of equality: $p = 0.58$).

Table 8: Technology Age Effect: Metropolitan vs. Micropolitan

	Metropolitan	Micropolitan
Modal Technology Age	0.124*** (0.032)	0.103*** (0.025)
Log Total Votes	-5.10*** (0.43)	-3.41*** (1.01)
Year FE	Yes	Yes
Observations	1,136	1,540
R^2	0.227	0.052

Standard errors clustered by CBSA. *** p<0.001.

A.3 Population-Weighted Results

Table 9 compares unweighted and population-weighted specifications. The technology coefficient attenuates under population weighting but remains significant.

Table 9: Population-Weighted Results

	(1)	(2)	(3)	(4)
	Unweighted	Pop-Weighted	Unw + 2008	Wgt + 2008
Modal Technology Age	0.075*** (0.016)	0.062*** (0.018)	0.052*** (0.012)	0.048*** (0.014)
GOP Share 2008 (%)			0.838*** (0.022)	0.812*** (0.028)
Log Total Votes	-4.71*** (0.28)	-3.82*** (0.35)	-0.68*** (0.18)	-0.52** (0.20)
Year FE	Yes	Yes	Yes	Yes
Weights	None	Votes	None	Votes
Observations	3,569	3,569	3,412	3,412
R^2	0.226	0.198	0.864	0.842

Standard errors clustered by CBSA. ** $p < 0.01$, *** $p < 0.001$.

A.4 Industry Controls

Table 10 shows results adding industry controls. The technology coefficient attenuates slightly but remains significant.

Table 10: Industry Structure Controls

	(1)	(2)
	+ Sector Count	+ Interaction
Modal Technology Age	0.068*** (0.016)	0.078*** (0.021)
N Industry Sectors	-0.025 (0.018)	-0.032 (0.025)
High Diversity		1.42 (1.52)
Tech Age \times High Div		-0.018 (0.028)
Year FE	Yes	Yes
Observations	3,569	3,569
R^2	0.228	0.228

Standard errors clustered by CBSA. *** $p < 0.001$.

B. Coefficient Stability (Oster Test)

Following [Oster \(2019\)](#), we assess selection on unobservables. Moving from a specification with only year fixed effects ($\beta = 0.134$, $R^2 = 0.05$) to one with full controls and 2008 baseline ($\beta = 0.052$, $R^2 = 0.86$), we compute $\delta^* = 2.8$. This exceeds the conventional threshold of 1.0, implying unobserved confounders would need to be nearly three times as important as observed controls to fully explain the estimated effect.

C. Pre-Trend Placebo Test

We regress the 2008–2012 change in GOP vote share on 2011 technology age. The coefficient is small (0.008, s.e. = 0.012) and statistically insignificant ($p = 0.51$). Technology did not predict pre-Trump voting changes, supporting our interpretation that the technology-voting relationship emerged specifically with Trump’s candidacy.

D. Data Appendix

D.1 Technology Vintage Data

The technology data were provided by Prof. Tarek Hassan (Boston University). For each CBSA-year, we observe approximately 45 industry-level modal age values and collapse to the CBSA-year level by computing the mean.

D.2 CBSA-County Crosswalk

We use the March 2020 CBSA delineation from the Census Bureau via NBER. After merging technology and election data, 896 CBSAs remain in the analysis sample.

D.3 Data Provenance and Correspondence

This research was initiated based on correspondence from Prof. Tarek Hassan (Boston University) to Prof. David Yanagizawa-Drott suggesting evaluation of the hypothesis that technological obsolescence is related to the rise of populist support.

Important Note on Data Correction: Earlier versions of this paper (APEP-0141, APEP-0143) incorrectly used simulated data rather than the actual technology vintage data provided by Prof. Hassan. Upon review, Prof. Hassan identified this error in February 2026:

“Hi David,

Thanks so much, this is great. I had a look and it seems it took data from some other (Acemoglu) paper, rather than from the dropbox link I sent. Can you try again with this link (this time as csv)?”

— Prof. Tarek Hassan, February 2026

This revision uses the correct data from Prof. Hassan’s Dropbox link (`modal_age.csv`). The results presented here reflect the actual technology vintage data as intended.

Data Sources:

- Technology vintage data: Provided by Prof. Tarek Hassan (Boston University) via Dropbox
- Election data (2008, 2012): MIT Election Data and Science Lab
- Election data (2016–2024): County-level compilations from replication files of [Bursztyn et al. \(2024\)](#)
- CBSA-county crosswalk: March 2020 Census Bureau delineation via NBER

E. Additional Figures




Figure 4: Distribution of Modal Technology Age Across U.S. Metropolitan Areas




Figure 5: Technology Age and Republican Vote Share: Binned Scatter Plot

Notes: Each point represents a ventile of CBSAs sorted by technology age.

Testing Causal vs. Sorting: Technology Age and Voting Gains

Tech age predicts gains from 2012–2016 (A) but not subsequent gains (B)




Figure 6: Levels vs. Gains: Testing for Causal Effects




Figure 7: 2016 Trump Vote Share vs 2008 McCain Vote Share by Technology Age Tercile

Notes: CBSAs above the diagonal gained GOP support from 2008 to 2016. Older-technology CBSAs (red) cluster above the line.