Do Voters Punish Inflation or Pay Cuts? Inflation and Real Wages in U.S. Elections

Online Appendix: Not for publication

Juan Felipe Riaño[a], Francesco Trebbi[b]

[a]*Department of Economics, Georgetown University, Washington, DC, USA*
[b]*Haas School of Business, University of California, Berkeley, Berkeley, CA, USA*

A B S T R A C T

This online appendix collects the material supporting *Do Voters Punish Inflation or Pay Cuts? Inflation and Real Wages in U.S. Elections*. It derives the paper's estimating equations from a discrete-choice model of voter behavior; documents the aggregate "missing intercept" time-series relationship between inflation and the incumbent-party vote from 1976 to 2024; and validates the Economic Policy Institute Family Budget Calculator against the Bureau of Labor Statistics Consumer Price Index and the Bureau of Economic Analysis Regional Price Parities. It then reports the full set of robustness exercises (alternative control sets, partisan pre-trend controls, nonlinearity, and log-odds specifications, together with a replication of Flavin's 2026 results), and presents the congressional-district estimates, an average-pay wage measure, additional summary statistics, and the county maps referenced in the main text.

Figures in the Appendix

ORCID(s): 0000-0003-4909-4127 (J.F. Riaño); 0009-0003-5599-4778 (F. Trebbi)

Tables in the Appendix

A. Microfoundation: From Voter Choice to Estimating Equations

This appendix derives the paper's estimating equations from a standard discrete-choice model of voting. The main text presents specifications (1)–(4) directly and discusses identification informally; here we show that these equations are the natural reduced-form implications of a logit model of two-party voter choice observed over two elections.

Individual voter choice

Consider voter i in county c during election t, who chooses between two parties $j \in \{D, R\}$. The voter's utility from party j is

$$U_{ijct} = V_{jct} + \varepsilon_{ijct}, \tag{1}$$

where V_{jct} is a systematic component common to all voters in county c and ε_{ijct} is an idiosyncratic preference shock. We decompose the systematic component as

$$V_{jct} = \alpha_{jc} + \delta_{jt} + \mathbf{X}'_{ct}\boldsymbol{\beta}_j, \tag{2}$$

where α_{jc} is a county–party fixed effect capturing the long-run partisan lean of county c toward party j (ideology, demographics, political culture); δ_{jt} is a national election–party shock (candidate effects, national mood, party platform); and $\mathbf{X}'_{ct}\boldsymbol{\beta}_j$ captures how local economic conditions affect support for party j.[1]

Aggregation to county-level vote shares

Under the logit assumption, the probability that voter i chooses party j takes the standard form:

$$\Pr(j \mid c, t) = \frac{\exp(V_{jct})}{\exp(V_{Dct}) + \exp(V_{Rct})}.$$

Invoking representativeness, the two-party vote share of party j in county c at election t is $s_{jct} = \Pr(j \mid c, t)$. Taking the log-odds of Republican relative to Democratic vote share and substituting (2):

$$\begin{aligned}\log\left(\frac{s_{Rct}}{s_{Dct}}\right) &= (V_{Rct} - V_{Dct}) \\ &= \underbrace{(\alpha_{Rc} - \alpha_{Dc})}_{\equiv \alpha_c} + \underbrace{(\delta_{Rt} - \delta_{Dt})}_{\equiv \delta_t} + \mathbf{X}'_{ct}\underbrace{(\boldsymbol{\beta}_R - \boldsymbol{\beta}_D)}_{\equiv \boldsymbol{\beta}} + \varepsilon_{ct},\end{aligned} \tag{3}$$

where ε_{ct} is a county-level unobservable that satisfies the conditional exogeneity restriction $E[\varepsilon_{ct} \mid \alpha_c, \delta_t, \mathbf{X}_{ct}] = 0$. The relative parameter $\boldsymbol{\beta} \equiv \boldsymbol{\beta}_R - \boldsymbol{\beta}_D$ captures the *differential* effect of economic conditions on Republican versus Democratic support: a positive coefficient means the economic variable favors Republicans on net.

First-differencing with two elections

We observe each county in exactly two presidential elections, $t \in \{2020, 2024\}$. Applying equation (3) to both years and taking the difference within county c:

$$\begin{aligned}\Delta \log\left(\frac{s_R}{s_D}\right)_c &\equiv \log\left(\frac{s_{Rc,2024}}{s_{Dc,2024}}\right) - \log\left(\frac{s_{Rc,2020}}{s_{Dc,2020}}\right) \\ &= \underbrace{(\delta_{2024} - \delta_{2020})}_{\kappa} + \Delta\mathbf{X}'_c\boldsymbol{\beta} + \Delta\varepsilon_c.\end{aligned} \tag{4}$$

The county fixed effects α_c are eliminated exactly. The national swing between elections collapses to a constant κ, which is absorbed by the regression intercept. Identification comes entirely from cross-county heterogeneity in the

[1] The model rests on four standard assumptions: (i) *additive separability* of systematic utility in county–party effects, national shocks, and local economic conditions; (ii) *logit errors*—the idiosyncratic shocks ε_{ijct} are i.i.d. Type-I extreme value (Gumbel) across voters, parties, and elections; (iii) *representativeness*—aggregate choice probabilities equal observed two-party vote shares; and (iv) *common coefficients*—the economic response parameters $\boldsymbol{\beta}_j$ are common across counties in the baseline specification.

changes $\Delta \mathbf{X}_c$, that is, from differences across counties in how local economic conditions evolved between 2020 and 2024. Any time-invariant county characteristic (baseline ideology, geography, demographics) is differenced out.

State-specific trends. Equation (3) assumes a single national swing δ_t common to all counties. In practice, electoral dynamics vary across states due to local political environments, down-ballot races, and turnout patterns. We therefore replace δ_t with state-specific swings $\delta_{s(c),t}$ in the level model. Upon first-differencing, these yield state fixed effects $\mu_{s(c)} \equiv \delta_{s(c),2024} - \delta_{s(c),2020}$, restricting identification to within-state variation in economic conditions.

Differential trends by baseline characteristics. Equation (4) includes only the *changes* $\Delta \mathbf{X}_c$ in time-varying covariates; any time-invariant county characteristic is absorbed by α_c and differenced out. However, we may want the electoral shift to vary systematically with observable baseline county characteristics. To do so, we augment the level model (3) with a term $\mathbf{Z}_c' \boldsymbol{\gamma} \cdot t$, where $\mathbf{Z}_c$ collects time-invariant (or slowly varying) characteristics such as urbanization share, population density, age structure, educational attainment, real GDP, and demographic composition. With two periods, first-differencing yields $\mathbf{Z}_c' \boldsymbol{\gamma}$: baseline county characteristics enter the differenced equation *in levels*. The coefficient vector $\boldsymbol{\gamma}$ captures *differential linear trends*—that is, systematic differences in the 2020-to-2024 electoral shift as a linear function of baseline county type. Crucially, including $\mathbf{Z}_c$ does not "control for" urbanization per se (which is time-invariant and already absorbed by α_c); rather, it allows the electoral *change* to differ across more urban versus more rural, more educated versus less educated counties, and so on.

Combining both extensions, the enriched first-differenced equation becomes:

$$\Delta \log\left(\frac{s_R}{s_D}\right)_c = \mu_{s(c)} + \Delta \mathbf{X}_c' \boldsymbol{\beta} + \mathbf{Z}_c' \boldsymbol{\gamma} + \Delta \varepsilon_c. \tag{5}$$

State fixed effects absorb state-level average swings; the level controls $\mathbf{Z}_c$ absorb differential linear trends across county types. Identification of $\boldsymbol{\beta}$ therefore relies on within-state, within-county-type variation in economic changes, that is, the residual variation in $\Delta \mathbf{X}_c$ after partialing out state means and linear functions of baseline characteristics.

From log-odds to linear vote-share changes

Equation (4) implies that the natural dependent variable is the change in log-odds of the Republican two-party vote share. In practice, we use the linear change in vote share Δs_c^R and the change in the Republican margin Δm_c as our primary outcome variables. This linear approximation is standard in the applied political economy literature (see, e.g., Autor, Dorn, Hanson and Majlesi, 2020; Charles, Hurst and Schwartz, 2019; Dippel, Gold, Heblich and Pinto, 2022) and is justified as follows.

Let $s \equiv s_R/(s_R + s_D)$ denote the Republican two-party share. The logistic transformation $\ell(s) = \log(s/(1-s))$ is approximately linear in s for values away from 0 and 1. A first-order Taylor expansion around the sample mean $\bar{s}$ gives

$$\Delta \log\left(\frac{s_R}{s_D}\right) \approx \frac{1}{\bar{s}(1-\bar{s})} \cdot \Delta s^R.$$

The coefficient $[\bar{s}(1-\bar{s})]^{-1}$ is a positive constant (approximately 4 when $\bar{s} \approx 0.5$, as in our sample). This is a scale factor that rescales coefficients but preserves their signs and relative magnitudes across specifications.

However, this approximation is exact only when two-party vote shares are far from 0 and 1 *and* when $\bar{s}(1-\bar{s})$ does not vary systematically across counties. In practice, counties with extreme partisan leans (very high or very low Republican share) have $\bar{s}(1-\bar{s})$ well below 0.25, which amplifies the log-odds change relative to the linear change.

Table A.1 verifies this empirically for total inflation using the paper's three specifications (baseline, gap-plus-price, and gap-only) under the most demanding controls. The two dependent variables, the standardized linear change Δs^R and the standardized log-odds change $\Delta \log(s_R/s_D)$, have a cross-county correlation of 0.93; three counties with 0% or 100% Republican two-party share are excluded from the log-odds specification. The qualitative conclusions are identical across both dependent variables: inflation is negative and significant in the baseline and gap-plus-price specifications; real wage growth is negative and significant in the gap-plus-price specification; and the gap alone is statistically insignificant when entered without prices. In the gap-plus-price specification, the log-odds coefficients are in fact larger in absolute value and more precisely estimated than their linear counterparts (e.g., the real-wage-growth coefficient is -0.136 ($t = -2.63$) under log-odds versus -0.052 ($t = -2.49$) under the linear DV), reflecting the logit model's natural weighting of vote-share variation.

Table A.1
Linear Vote-Share Change vs. Log-Odds Change as Dependent Variable (Total Inflation)

	Baseline		Gap + Price		Gap Only	
	(1) Δs^R	(2) $\Delta \log \frac{s_R}{s_D}$	(3) Δs^R	(4) $\Delta \log \frac{s_R}{s_D}$	(5) Δs^R	(6) $\Delta \log \frac{s_R}{s_D}$
Inflation	−0.083***	−0.045**	−0.121***	−0.143***		
	(0.024)	(0.022)	(0.027)	(0.044)		
Nominal Wage Growth	−0.037**	−0.098***				
	(0.015)	(0.037)				
Real Wage Growth			−0.052**	−0.136***	0.003	−0.071
			(0.021)	(0.052)	(0.026)	(0.049)
R^2	0.413	0.430	0.413	0.430	0.408	0.423
Observations	3,102	3,099	3,102	3,099	3,102	3,099
Controls	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Observations are at the county level. All regressions use total EPI Family Budget inflation (2021–2024) and the most demanding specification (full controls and state fixed effects). Odd columns use the standardized change in two-party Republican vote share (Δs^R) as the dependent variable; even columns use the standardized change in log-odds ($\Delta \log(s_R/s_D)$). Columns (1)–(2) correspond to the baseline specification (equation 1); (3)–(4) to the gap-plus-price reparameterization (equation 3); (5)–(6) to the restricted gap-only model (equation 4). The two dependent variables have a cross-county correlation of 0.93; 3 counties with 0% or 100% Republican two-party share are dropped in the log-odds specification. All variables standardized. Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

The vote margin $\Delta m_c = \Delta(v^R - v^D)$ is an affine transformation of Δs^R and inherits the same properties. Turnout, the third outcome, is not derived from the two-party logit but is motivated by the same accountability framework: voters who experience adverse (or favorable) economic conditions may increase (or decrease) their propensity to participate, a channel that maps into a reduced-form specification with the same right-hand-side structure.

Deriving the paper's estimating equations

We now apply equation (5). The time-varying economic conditions $\mathbf{X}_{ct}$ comprise local inflation $P_c^{(k)}$ for expenditure category k and nominal wage growth W_c, so $\Delta \mathbf{X}_c$ includes $\Delta P_c^{(k)}$ and ΔW_c. The baseline characteristics $\mathbf{Z}_c$ and state fixed effects $\mu_{s(c)}$ are as defined above. The estimating equation is therefore:

$$\Delta y_c = \beta_1 \Delta P_c^{(k)} + \beta_2 \Delta W_c + \mathbf{Z}'_c \boldsymbol{\gamma} + \mu_{s(c)} + \varepsilon_c, \tag{1}$$

which is the baseline specification in the main text. Here β_1 captures how Republican electoral gains covary with local price growth holding nominal wages fixed, and β_2 captures the association with nominal wage growth holding prices fixed.

Defining the category-specific wage–price gap $G_c^{(k)} \equiv \Delta W_c - \Delta P_c^{(k)}$ and substituting $\Delta W_c = G_c^{(k)} + \Delta P_c^{(k)}$ into the baseline yields:

$$\Delta y_c = \delta_1 G_c^{(k)} + \delta_2 \Delta P_c^{(k)} + \mathbf{Z}'_c \boldsymbol{\gamma} + \mu_{s(c)} + \varepsilon_c, \tag{3}$$

where $\delta_2 = \beta_1 + \beta_2$ and δ_1 equals β_2 up to a positive rescaling factor determined by the standardization convention: because $G_c^{(k)}$, ΔW_c, and $\Delta P_c^{(k)}$ are each standardized to zero mean and unit variance, $\delta_1 = \beta_2 \sqrt{2(1 - \rho_{PW})}$, where ρ_{PW} is the partial correlation between standardized wage and price growth net of state fixed effects and baseline controls. Because this rescaling is a positive scalar, the sign of δ_1, its t-statistic, and the p-value for $H_0^{(1)} : \delta_1 = 0$ coincide with those of β_2. This is an algebraic rotation: it leaves fitted values and R^2 unchanged. Its value is interpretive. The coefficient δ_1 measures how outcomes covary with real wage growth (the purchasing-power gap), while δ_2 captures the residual association of nominal price growth *conditional on* the gap. The pocketbook voting restriction $H_0^{(1)}$ ($\delta_1 = 0$) tests whether real wage growth is electorally relevant; the pure pocketbook restriction $H_0^{(2)}$ ($\delta_2 = 0$) holds if and only if $\beta_1 = -\beta_2$, meaning voters respond to inflation only through its effect on real wages, with no residual price salience.

Finally, the gap-only specification

$$\Delta y_c = \delta_1 G_c^{(k)} + \mathbf{Z}_c' \boldsymbol{\gamma} + \mu_{s(c)} + \varepsilon_c \tag{4}$$

imposes this restriction and tests whether real wage growth alone captures the relevant covariation.

B. Aggregate Time-Series Evidence and the Missing Intercept

Appendix A derives the estimating equations and shows that identification of $\boldsymbol{\beta}$ comes from residual within-state, within-county-type, cross-county variation in $\Delta \mathbf{X}_c$, after the common national swing, state-specific linear trends, and differential trends across county types have been absorbed. As emphasized by Wolf (2023), residual cross-sectional estimates of this form recover a relative response and do not, in general, identify the aggregate general-equilibrium response of the national economy to an aggregate shock, nor the response of any specific county type. The absorbed object—the "missing intercept"—can differ from the within-state cross-section in both magnitude and sign whenever the aggregate shock activates general-equilibrium channels that are absent from the local cross-section.

Figure B.1 visualizes this absorbed object in the longest time-series analog available for U.S. presidential elections. For each election from 1976 to 2024, the horizontal axis plots four-year term-average CPI-U All Items inflation and the vertical axis plots the change in the incumbent party's two-party popular-vote share relative to the previous election. The ordinary-least-squares fit through the thirteen points is negative, with a Pearson correlation near −0.49: elections that follow higher-inflation terms tend to be ones in which the incumbent party loses ground. Each observation summarizes national conditions and the national election outcome, and the variation across observations is driven by the common national, state-trend, and county-type-trend components that our specification projects out. A within-state, within-county-type cross-sectional estimate and an aggregate time-series estimate therefore identify different objects. The figure documents what the paper's design does not estimate.

Figure B.1: Aggregate inflation and change in the incumbent party's two-party popular-vote share, U.S. presidential elections 1976–2024

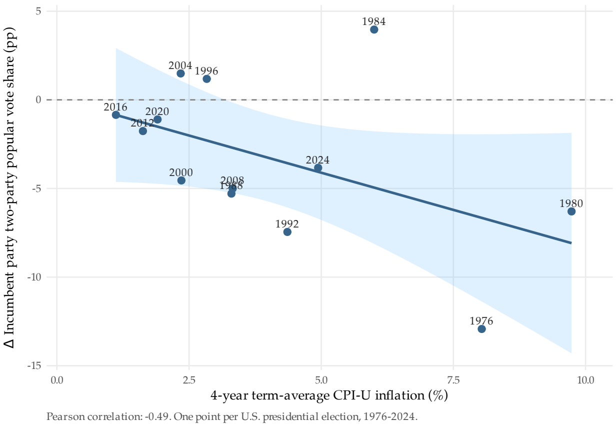

Notes: One observation per election. Term inflation is the simple average of year-over-year growth in the annual CPI-U All Items index over the four years ending in the election year. Δ incumbent share is the difference between the incumbent party's two-party popular-vote share at election t and the same party's two-party share at $t-4$; 1972 shares are taken from Federal Election Commission totals. Shaded band shows the 95% confidence interval of the ordinary-least-squares fit.

C. Validation of the EPI Family Budget Measure

C.1. Concerns and caveats

Several concerns may arise when using the EPI Family Budget Calculator (FBC) changes to proxy for local inflation. The FBC is produced by the Economic Policy Institute, a labor-aligned research organization, rather than by a federal statistical agency. It measures the cost of a normatively-defined "modest yet adequate" standard of living rather than the cost of a fixed market basket, and its components (Department of Housing and Urban Development

(HUD) Fair Market Rents for housing, the U.S. Department of Agriculture (USDA) Low-Cost Food Plan for food, state child-care market-rate surveys for childcare, and healthcare premiums from the Medical Expenditure Panel Survey (MEPS)) are revised on independent schedules. The three FBC releases used in this paper (the 2022, 2024, and 2025 releases, corresponding to reference years 2021, 2023, and 2024) are issued under the same institutional framework but are not documented as methodologically identical across vintages.

EPI is explicit on this point. The FBC documentation states that “much of the source data used to construct the Family Budget Calculator has been changed over the years, and the EPI methodology for measuring costs has often had to change as well. Further, many people—particularly researchers—want to use the Family Budget Calculator as a price index over time. This is not an appropriate use—the Family Budget Calculator does not just measure prices, it also measures spending. The source data for the rent component, for example, are likely measuring the cost of apartments that change in size and quality over time. A similar issue holds true for health care—if any of the rise in health care costs constitutes quality improvements, then this cannot be used as a pure price index” (Economic Policy Institute, 2024). EPI recommends instead that researchers “use the Family Budget Calculator as a rough proxy for price differences across areas and then anchor these to changes in region-specific price indexes from other sources if one wanted measures of price level changes over time across areas.”

The identification strategy in this paper follows that recommendation. FBC growth rates are never treated as a stand-alone price index. The dependent variables in every specification are differences over time, $\Delta y_c = y_{c,2024} - y_{c,2020}$, which by construction eliminate all time-invariant county characteristics. State fixed effects on these differences absorb state-specific linear trends in the underlying levels, including any state-common (or higher-level) methodology revision or shift in EPI's reference standards. EPI's reference standards are set at national or broad-area levels: USDA food plans are national, HUD Fair Market Rents are metropolitan-wide, and MEPS premium data are state-level. The “standards” component that EPI cautions against is therefore approximately common across counties within a state and is absorbed by the fixed effects. The formal decomposition in Section 8 shows that the remaining identifying variation isolates the residual price component of FBC growth. The remainder of this appendix assesses the spatial price signal in the FBC against two independent external benchmarks. Section C.2 compares FBC inflation to the Bureau of Labor Statistics (BLS) Consumer Price Index for All Urban Consumers (CPI-U) across the self-representing metropolitan Primary Sampling Units (PSUs). Section C.3 compares both the level and growth of FBC costs to the Bureau of Economic Analysis (BEA) Regional Price Parities (RPPs) across all 381 U.S. Metropolitan Statistical Areas (MSAs). Section C.4 summarizes the results.

C.2. Comparison with the CPI-U

The CPI-U is the most widely used U.S. price measure. It is constructed from price quotes collected in 75 PSUs, of which 23 are “self-representing” metropolitan areas large enough that BLS publishes individual CPI indexes. Figure 2(a) maps both types of PSUs. We restrict attention to the 21 self-representing PSUs in the continental United States (excluding Anchorage and Honolulu), computing CPI inflation as the percentage change in the annual average CPI-U index between 2021 and 2024 (not seasonally adjusted). County-level EPI Family Budget inflation is aggregated to the PSU level using 2020 Census county population as weights.

The EPI Family Budget Calculator includes seven expenditure categories: food, housing, transportation, healthcare, childcare, taxes, and other necessities. Four of these (food, housing, transportation, and healthcare) have direct counterparts in the CPI. The remaining three (childcare, taxes, and other necessities) have no metropolitan-level CPI equivalent but together account for over a third of the family budget. Table C.1 reports the mean and median annual expenditure and budget share for each category.

We construct a restricted EPI basket containing only the four CPI-comparable categories and compare its inflation to CPI-U All Items. Figure C.1 plots the result. The two measures are strongly correlated ($r = 0.74$). Metropolitan areas where CPI inflation was highest (Miami, Tampa, Phoenix) exhibit the highest EPI inflation in CPI-comparable categories, and vice versa.

The CPI does not cover childcare or taxes in its metropolitan indexes, yet both exhibit substantial geographic variation relevant for household economic circumstances. The CPI also provides no county-specific estimates for the rural United States: individual indexes exist for only 23 self-representing PSUs, and the remainder of the country is covered through four regional aggregates. The BEA RPPs, examined next, publish indexes for all 381 MSAs and for the nonmetropolitan residual of every state.

Table C.1
EPI Family Budget Composition, 2021

Category	Mean ($)	Median ($)	Share (%)	CPI counterpart
Food	9,243	9,115	11.8	✓
Housing	10,177	8,988	13.0	✓
Transportation	1,232	1,229	1.6	✓
Healthcare	16,081	15,774	20.6	✓
Childcare	12,841	12,165	16.4	
Taxes	8,028	7,635	10.3	
Other necessities	7,037	6,596	9.0	
Total	78,192	76,112	100.0	

Notes: Observations are at the county level. Annual expenditures for a two-parent, two-child family from the EPI Family Budget Calculator (2021 values). Mean and median computed across all counties in the regression sample. The "CPI counterpart" column indicates categories with a direct equivalent in the CPI-U index at the metropolitan level.

C.3. Comparison with the BEA RPPs

The RPPs are geographic price indexes measuring cost-of-living differences across U.S. metropolitan and nonmetropolitan areas (Bureau of Economic Analysis, 2025). They are constructed from BLS CPI price data, American Community Survey (ACS) housing-rent microdata, and BEA Personal Consumption Expenditure (PCE) weights, aggregated via a multilateral Geary index. The published RPP is normalized so that the U.S. average equals 100 each year. We compare the FBC to the RPPs at two horizons: the level of cost-of-living differences across areas in a single year, and the growth of area-level prices over 2021–2024. The growth comparison uses BEA's Implicit Regional Price Deflator (MAIRPD), defined as the product of the RPP and the national PCE price index; this is BEA's prescribed formula for converting the RPP into an absolute regional inflation rate (Bureau of Economic Analysis, 2025, §4.31). County-level EPI dollar costs are aggregated to the MSA boundaries BEA uses (the March 2020 Office of Management and Budget (OMB) delineation) with 2024 Census population weights; counties outside any Core-Based Statistical Area (CBSA) are assigned to their state's nonmetropolitan residual.

Levels. Figure C.2 plots the EPI FBC total family budget (in thousands of dollars, national average $97,000 in 2024) against BEA's all-items RPP (U.S. = 100) for all 381 MSAs, in 2021 (panel a) and 2024 (panel b). The population-weighted Pearson correlation is $\rho = 0.77$ in 2021 and $\rho = 0.80$ in 2024. The high-cost metros (San Francisco, San Jose, Oxnard) and the low-cost metros (Eagle Pass, McAllen, Brownsville) are the same in both measures, and the correlation is stable across the two FBC releases that bracket the analysis window.

Matched-basket growth comparison. Three features of the RPP methodology inform the construction of the growth comparison. First, the RPPs "substitute…national price levels" for categories whose regional CPI estimates are excessively volatile, including hospital services, physician services, prescription drugs, college and primary-school tuition, and childcare (Bureau of Economic Analysis, 2025, §4.11); the RPPs therefore carry no regional variation in childcare, most medical services, or education. Second, rural counties are assigned the price level of "the CPI regional index area in which they are located" (Bureau of Economic Analysis, 2025, §4.12), so the RPP has limited cross-county signal outside of metropolitan areas. Third, BEA notes that the underlying CPI "was not designed for place-to-place comparisons, which can lead to volatility in the results, especially across areas that have few observations or where sampling sizes are unevenly distributed" (Bureau of Economic Analysis, 2025, §2.2).

The comparison therefore uses matched EPI baskets and a sample restricted to metros where the underlying CPI sample is densest. The *BEA-comparable basket* consists of food, housing, transportation, healthcare, and other necessities, which together correspond to the five published RPP sub-aggregates; childcare and taxes are dropped because neither has a PCE counterpart. The *BEA-strict basket* additionally drops healthcare, because BEA imposes national price levels on its largest healthcare subcomponents. The *housing-only* comparison isolates the single category in which EPI (using HUD Fair Market Rent at the county level) and BEA (using ACS tenant-rent microdata at the Public Use Microdata Area (PUMA) level) both rest on directly comparable rental data. Correlations are computed on the 100 most populous MSAs, weighted by population.

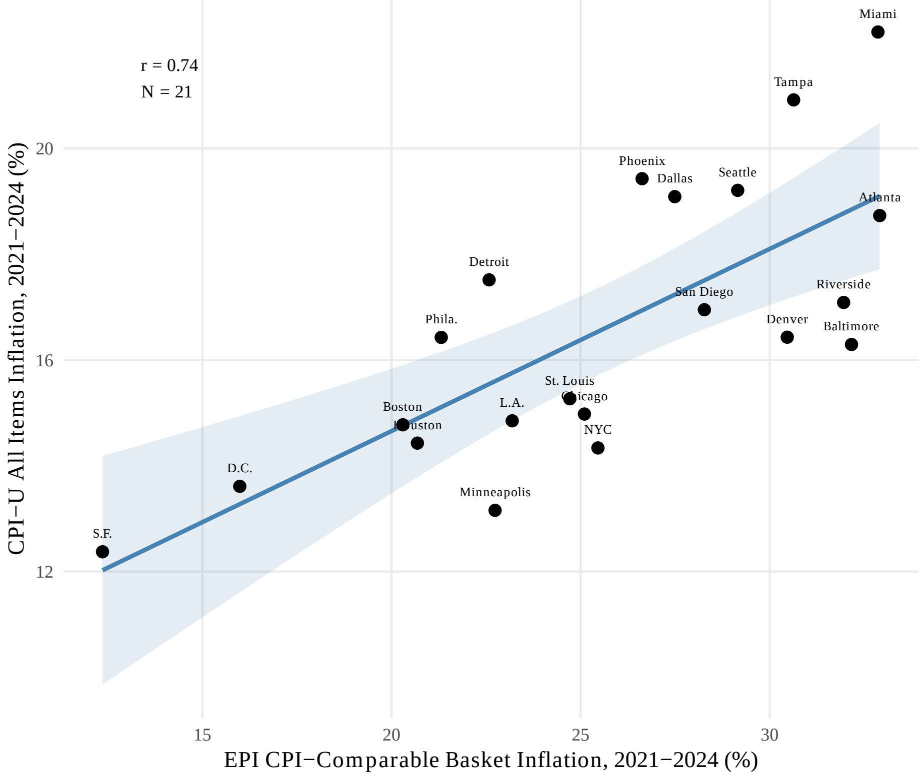

Figure C.1: EPI CPI-Comparable Basket Inflation vs. CPI-U All Items Inflation, 2021–2024

Notes: Each point represents one of 21 self-representing CPI Primary Sampling Units in the continental United States. The horizontal axis plots population-weighted average EPI Family Budget inflation computed using only the four expenditure categories with direct CPI counterparts: food, housing, transportation, and healthcare (excluding taxes, childcare, and other necessities). The vertical axis plots official CPI-U All Items inflation (annual average, not seasonally adjusted). The solid line is the OLS fit. CPI data retrieved from the Bureau of Labor Statistics public API.

Table C.2 summarizes the results. On the BEA-strict basket and the top-100 metro sample, the population-weighted Pearson correlation between EPI and BEA inflation over 2021–2024 is $\rho = 0.51$. On the housing-only comparison, $\rho = 0.55$ in the top 100 MSAs and $\rho = 0.66$ in the top 25. The pattern across rows is monotone: tightening the sample to the metros BEA measures more precisely raises the correlation (the right-most columns rise from 0.42 at all 381 metros to 0.54–0.66 in the top 25), and tightening the basket to match what BEA prices regionally also raises it (the BEA-strict column is uniformly higher than the BEA-comparable column). Figure C.3 plots the headline scatters: panel (a) the BEA-strict basket comparison, panel (b) the housing-only comparison.

The residual gap between EPI and BEA concentrates in categories where the RPPs carry no regional variation: childcare, the services side of healthcare, and the rural nonmetropolitan United States. EPI's primary sources for these categories (state-level child-care market-rate surveys, MEPS state-level premium data, and HUD county-level Fair Market Rents) are purpose-built for place-to-place comparison. The basket weights also differ: EPI's family-budget weights place roughly 40% on housing and 10–15% on childcare, while BEA's PCE weights place roughly 15% on housing and carry a near-zero regional share on childcare. The EPI FBC therefore measures the portion of the area-level cost-of-living signal that the RPPs also measure, together with additional variation that the RPPs, by construction, do not.

C.4. Summary

Across the two benchmarks, EPI Family Budget inflation aligns with established federal price measures on the dimensions where alignment is expected. Against the CPI-U, the four-category overlap basket yields $r = 0.74$ over the 21 self-representing metropolitan PSUs. Against the BEA RPPs, the cross-sectional level of cost-of-living registers

Figure C.2: EPI Family Budget Total vs. BEA Regional Price Parity (all items), 2021 and 2024

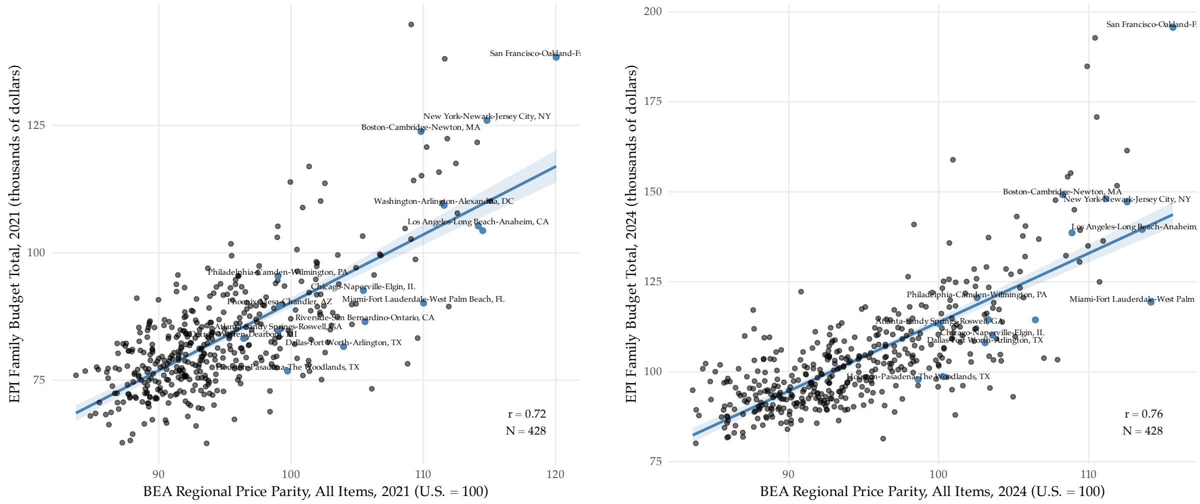

(a) 2021 **(b) 2024**

Notes: Each point is one of 381 U.S. metropolitan statistical areas (plus state nonmetropolitan residuals). The horizontal axis is the BEA all-items Regional Price Parity (U.S. = 100); the vertical axis is the EPI Family Budget total cost for a two-parent, two-child family, in thousands of current dollars. The blue line is the OLS fit and the shaded band is its 95% confidence interval. The 15 most populous MSAs are shown in steelblue with labels. Population-weighted and simple Pearson correlations are printed in the lower-right; $N = 428$ in both panels.

Table C.2

Population-weighted Pearson correlations of EPI Family Budget inflation with BEA Implicit Regional Price Deflator inflation, 2021–2024

Sample	N	EPI Housing only	BEA-comparable basket	BEA-strict basket
All metros + state-nonmetro	428	0.390	0.332	0.416
All metros (MARPP)	381	0.404	0.350	0.440
Top 200 metros by population	200	0.478	0.377	0.465
Top 100 metros by population	100	0.550	0.413	0.511
Top 50 metros by population	50	0.628	0.397	0.536
Top 25 metros by population	25	0.659	0.404	0.560

Notes: Each row reports the population-weighted Pearson correlation between three EPI inflation measures and BEA's Implicit Regional Price Deflator inflation over 2021–2024, on progressively tighter samples of U.S. metropolitan statistical areas. *EPI Housing only* uses the EPI housing-inflation series, compared to BEA's published housing-component RPP growth. *BEA-comparable basket* drops childcare and taxes from the EPI total (neither has a BEA PCE counterpart). *BEA-strict basket* additionally drops healthcare, because BEA substitutes national price levels for hospital, physician, and prescription-drug services (Bureau of Economic Analysis, 2025, §4.11). The sample cascade addresses BEA's own caveat that the underlying CPI is noisier for metros with smaller price-survey samples (Bureau of Economic Analysis, 2025, §2.2). Weights are 2024 county population summed to the MSA.

$\rho = 0.80$ in 2024 and $\rho = 0.77$ in 2021 across all 381 MSAs; the matched-basket inflation correlation reaches $\rho = 0.51$ on the BEA-strict basket and $\rho = 0.55$–0.66 on housing alone in the 100 most populous MSAs. The agreement of two benchmarks built with distinct methodologies is consistent with the FBC capturing a genuine cross-county price signal.

The exercise also clarifies what the FBC adds relative to each benchmark. It prices childcare and state-level taxes at sub-state granularity, which the CPI does not. It retains regional variation in categories that the BEA RPP averages away by design—childcare and the high-importance services side of healthcare (hospital, physician, and prescription-drug services)—and provides county-level estimates for the rural United States, where the RPP collapses to four regional

Figure C.3: EPI and BEA inflation, 2021–2024, top 100 MSAs

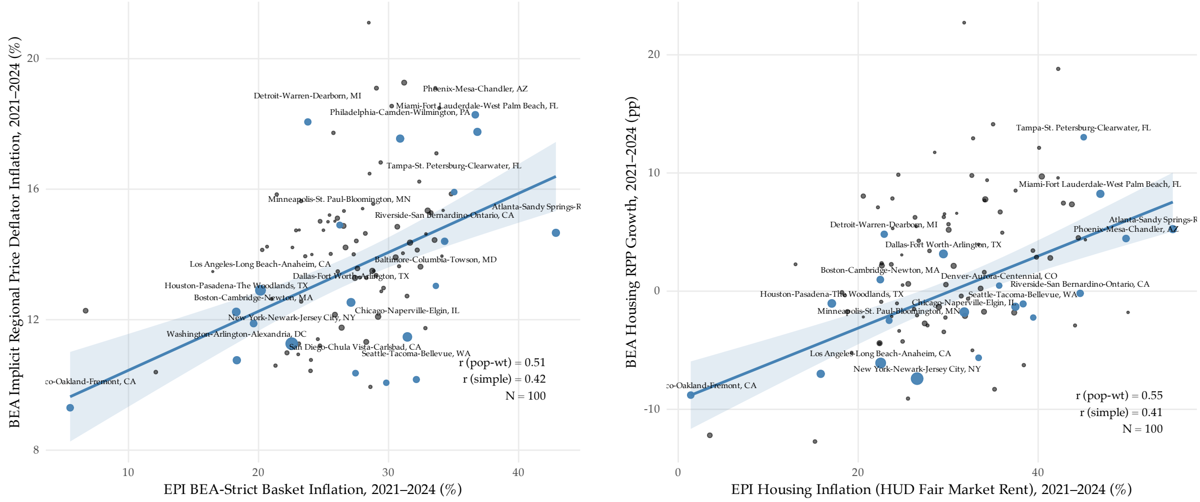

(a) BEA-strict basket **(b) Housing only**

Notes: Each point is one of the 100 most populous U.S. metropolitan statistical areas, sized by 2024 county population summed to the MSA. Panel (a) plots EPI Family Budget inflation 2021–2024 on the BEA-strict basket (food + housing + transportation + other necessities) against BEA's Implicit Regional Price Deflator inflation over the same window. Panel (b) restricts the comparison to the housing component in both sources (HUD Fair Market Rent for EPI; ACS-rent-based RPP for BEA). The 20 most populous MSAs are highlighted in steelblue with labels. The blue line is the OLS fit weighted by population and its 95% confidence band; the lower-right annotation reports both population-weighted and simple Pearson correlations.

aggregates. These categories and geographies are those most relevant for a study of county-level electoral responses to local cost-of-living changes.

Consistent with EPI's own caution against using the FBC as a pure time-series price index, identification exploits within-state, cross-county variation in changes. First-differencing the dependent variable removes time-invariant county characteristics, and state fixed effects on the differences absorb state-specific linear trends in levels, including any methodology revision or reference-standard shift common at the state level or above.

A design that exploits county-level variation in cost-of-living changes over a three-year horizon requires four features not jointly available from published federal series: county-level granularity (the CPI publishes individual indexes for 23 PSUs and the RPP for roughly 430 areas); rural-county coverage (both the CPI and the RPP collapse rural areas into broad regional aggregates); a basket that includes childcare, state-level taxes, and rental housing at the Fair Market Rent standard; and inflation measured over a cumulative horizon longer than a single year. The EPI Family Budget Calculator supplies all four.

D. Robustness to Housing-Price and Unemployment Controls

Related work on cross-sectional electoral responses to inflation, notably Flavin (2026), includes county-level house-price growth and contemporaneous unemployment among the control variables, on the intuition that these capture broader local economic conditions. An analyst concerned that the inflation-vote relationship is confounded with such conditions has two distinct options for augmenting the regression, and the distinction between them matters for our setting. The first option is a *pre-shock* snapshot of local economic conditions, measured before the inflation episode under study—for example, the December 2020 level of the Zillow Home Value Index or the 2020 BLS Local Area Unemployment Statistics (LAUS) unemployment rate. These are proper pre-treatment covariates, and if they carry information not already contained in the baseline covariate set their inclusion would tighten the estimates without distorting them. The second option is a *contemporaneous* measure of the same variables, such as the USDA Economic Research Service (ERS) unemployment rate in 2022 or 2023 or the FHFA House Price Index annual changes through 2024. These are post-treatment outcomes of the same inflation–wage shock the paper studies: the pandemic-era demand surge that lifted local prices also tightened local labor markets, and roughly forty percent of the EPI Family Budget

Table C.3
Summary Statistics — County Level, Supplementary Variables

Variable	Mean	SD	Min	P25	P50	P75	Max	N
Panel A: Electoral Outcomes — Levels								
Two-party Rep. vote share 2020 (%)	66.1992	16.2658	5.5330	57.0047	69.5876	78.7896	96.8864	3102
Two-party Rep. vote share 2024 (%)	67.9132	15.8099	6.6853	59.0949	71.3618	80.2003	100.0000	3100
Turnout 2020 (%)	63.7016	9.8037	3.3900	57.3929	63.8743	70.3684	100.0000	3102
Turnout 2024 (%)	63.2457	11.0023	0.0000	56.2957	63.3827	70.7273	100.0000	3102
Rep. vote margin 2020 (pp)	31.8912	32.0096	-86.7524	13.7477	38.5096	56.6332	93.0909	3102
Rep. vote margin 2024 (pp)	35.3622	31.1832	-83.1600	17.9302	42.1083	59.6479	98.6502	3102
Panel B: Family Budget Costs — Levels								
FBC food budget 2021 (thousands)	9.2391	0.9945	6.4162	8.6349	9.1146	9.6542	20.8676	3102
FBC food budget 2024 (thousands)	11.6642	1.2019	8.9012	10.8575	11.4770	12.1617	25.7254	3102
FBC housing budget 2021 (thousands)	10.1618	3.1427	5.7813	8.6160	8.9760	10.5964	48.0286	3102
FBC housing budget 2024 (thousands)	12.7727	4.0376	6.4494	10.5240	11.1360	13.3680	48.6480	3102
FBC childcare budget 2021 (thousands)	12.8288	4.0993	4.0081	10.2217	12.1601	15.0507	42.0460	3102
FBC childcare budget 2024 (thousands)	15.1232	5.2395	4.6152	11.8018	14.2358	17.4764	52.9310	3102
FBC transportation budget 2021 (thousands)	1.2320	0.0822	0.2215	1.1816	1.2289	1.2829	1.5568	3102
FBC transportation budget 2024 (thousands)	1.6188	0.1417	0.5311	1.5272	1.6072	1.7133	2.1743	3102
FBC healthcare budget 2021 (thousands)	16.0818	2.6855	10.4428	14.3373	15.7743	17.3152	27.9849	3102
FBC healthcare budget 2024 (thousands)	18.3426	3.1734	12.1059	16.2535	17.7358	19.6998	30.8276	3102
FBC other necessities budget 2021 (thousands)	7.0297	1.3713	5.0715	6.2767	6.5942	7.2591	21.7264	3102
FBC other necessities budget 2024 (thousands)	8.2219	1.6266	5.6036	7.2661	7.6528	8.5265	21.5459	3102
FBC taxes budget 2021 (thousands)	8.0140	2.7168	3.1767	6.3354	7.6307	9.1552	33.8597	3102
FBC taxes budget 2024 (thousands)	11.6300	3.7425	6.2377	9.3794	10.9377	12.6431	47.6226	3102
FBC total budget 2021 (thousands)	78.1394	10.8036	59.6852	71.1866	76.0774	82.6065	172.8881	3102
FBC total budget 2024 (thousands)	97.1806	14.2348	75.0669	88.6139	94.1194	102.0835	222.7196	3102
Panel C: Wage Levels								
Annual total wages, 2021 (billions USD, CD agg.)	2.9541	13.8108	0.0026	0.1275	0.3649	1.1812	349.8209	3102
Annual total wages, 2024 (billions USD, CD agg.)	3.5142	16.1560	0.0023	0.1544	0.4370	1.4210	407.2340	3102
Annual average pay, 2021 (thousands USD, CD agg.)	47.1259	11.0836	26.7440	40.6890	44.9710	50.7300	185.2530	3102
Annual average pay, 2024 (thousands USD, CD agg.)	53.8720	12.0381	31.8590	46.7360	51.4415	57.9110	197.0430	3102
Panel D: Real Wage Growth (Average Pay – Inflation)								
Real wage growth: total (Avg Pay)	-9.8354	10.1513	-77.5309	-15.4514	-9.6027	-3.6100	80.9735	3102
Real wage growth: food (Avg Pay)	-11.9140	10.8512	-71.5596	-18.5245	-11.7965	-5.4017	63.0954	3102
Real wage growth: transportation (Avg Pay)	-16.7189	9.1303	-149.7011	-21.2437	-16.5210	-12.0456	69.5834	3102
Real wage growth: housing (Avg Pay)	-11.0145	11.8773	-78.7609	-17.7476	-10.5013	-4.1554	84.1568	3102
Real wage growth: childcare (Avg Pay)	-5.9893	30.5612	-231.5346	-20.5612	-0.2569	13.5161	104.9722	3102
Real wage growth: healthcare (Avg Pay)	-0.1442	16.1322	-102.9763	-8.5698	0.4594	9.5742	85.8988	3102
Real wage growth: othernecessities (Avg Pay)	-2.3352	8.9276	-57.6257	-7.3355	-2.6723	2.3682	82.9040	3102
Real wage growth: taxes (Avg Pay)	-32.9301	21.3302	-134.7422	-44.2879	-30.4289	-18.8346	69.2668	3102

Notes: Observations are U.S. counties in the county estimation sample ($N = 3{,}102$). Panel A reports 2020–2024 changes in county-level presidential vote shares, vote margins, and turnout, computed from MIT Election Lab county returns. Panel B reports 2021–2024 percent changes in the Economic Policy Institute Family Budget Calculator costs, two-parent two-child family, by expenditure category. Panel C reports 2021–2024 percent changes in total covered wages and in average annual pay, both from the Quarterly Census of Employment and Wages (QCEW). Panel D reports the category-specific wage–price gap, defined as nominal wage growth minus FBC inflation. Panel E reports the baseline covariate set used in every specification, drawn from the 2020 Census and the American Community Survey 2019–2023.

Calculator basket is the HUD Fair Market Rent, which co-moves with the FHFA index by construction. Conditioning on such contemporaneous outcomes is a textbook case of conditioning on a bad control in the sense of Angrist and Pischke (2009): the included regressor sits downstream of the treatment, and the resulting coefficient is biased relative to the reduced-form association of interest.

On these grounds, the specification reported in the main text omits both types of controls: the pre-shock versions because they add little beyond the baseline covariate set, and the contemporaneous versions because they are endogenous to the shock. The remainder of this appendix reports the exercise in full, separating the two cases.

Table D.1 presents four gap-reparameterized specifications for the change in the Republican two-party vote share. Column (1) reproduces the preferred baseline from Table 3. Column (2) adds pre-shock levels: the December 2020 Zillow Home Value Index and the 2020 LAUS unemployment rate. Column (3) adds the contemporaneous regressors used in the main joint specification of Flavin (2026): ERS unemployment rates for 2022 and 2023. Column (4) adds the union of the two sets plus the four annual FHFA House Price Index changes for 2021–2024. "Real Wage Growth" denotes the standardized wage–price gap $G_c^{(k)}$, and "Inflation" the residual inflation coefficient δ_2 from equation (3). Heteroskedasticity-robust standard errors are reported in parentheses and state-clustered bootstrap standard errors (100 replications) in brackets, matching Table 3. Table D.2 repeats the exercise for the turnout outcome.

The pre-shock specification in column (2) is business as usual. The real-wage-growth coefficient moves from -0.052 to -0.061 (both significant at the five-percent level under robust inference); the residual inflation coefficient moves from -0.121 to -0.112 (both significant at the one-percent level under both inference procedures). Signs, relative magnitudes, and statistical significance are preserved, and the two sets of estimates are well within each other's

Table C.4
Summary Statistics — Congressional District Level, Supplementary Variables

Variable	Mean	SD	Min	P25	P50	P75	Max	N
Panel A: Electoral Outcomes — Levels								
Two-party Rep. vote share 2020 (%)	48.1657	20.8874	0.0000	34.6059	48.9706	62.5445	100.0000	1270
Two-party Rep. vote share 2024 (%)	50.9382	21.9529	0.0000	36.5745	50.6321	65.0498	100.0000	1270
Turnout 2020 (%)	63.1287	17.6396	-0.0004	52.8163	61.4231	71.2092	100.0000	1257
Turnout 2024 (%)	61.9308	18.3920	0.0001	50.8241	60.6204	71.1548	100.0000	1257
Rep. vote margin 2020 (pp)	-3.3874	39.1407	-100.0000	-29.8500	-2.0385	24.7131	100.0000	1270
Rep. vote margin 2024 (pp)	1.8185	40.3669	-100.0000	-26.8510	1.2643	29.7037	100.0000	1270
Panel B: Family Budget Costs — Levels								
FBC food budget 2021 (CD agg) (thousands)	9.8375	1.1600	7.1188	9.0619	9.6349	10.3740	18.5573	1270
FBC food budget 2024 (CD agg) (thousands)	12.5810	1.1980	10.0737	11.7378	12.5490	13.3028	19.4313	1270
FBC housing budget 2021 (CD agg) (thousands)	15.0642	5.8155	8.1068	10.5509	13.5426	17.4136	45.8120	1270
FBC housing budget 2024 (CD agg) (thousands)	19.3337	6.9091	10.3021	13.7275	17.4607	23.3714	42.9906	1270
FBC childcare budget 2021 (CD agg) (thousands)	16.7284	5.6626	5.7473	13.1056	15.8443	18.8429	39.2537	1270
FBC childcare budget 2024 (CD agg) (thousands)	20.7614	7.2021	8.6336	15.7421	18.8676	25.0457	52.9310	1270
FBC transportation budget 2021 (CD agg) (thousands)	1.1685	0.1336	0.2216	1.1454	1.1875	1.2237	1.4320	1270
FBC transportation budget 2024 (CD agg) (thousands)	1.4822	0.1545	0.5312	1.4390	1.5089	1.5673	1.7434	1270
FBC healthcare budget 2021 (CD agg) (thousands)	14.6680	2.0951	10.4439	13.1782	14.4231	15.8873	25.8442	1270
FBC healthcare budget 2024 (CD agg) (thousands)	17.1557	2.6938	12.1059	15.1444	16.8358	18.2754	29.2819	1270
FBC other necessities budget 2021 (CD agg) (thousands)	9.0229	2.4356	5.8688	7.1612	8.4246	10.0513	21.1467	1270
FBC other necessities budget 2024 (CD agg) (thousands)	10.7379	2.6181	7.0613	8.6705	10.0773	12.3584	19.6939	1270
FBC taxes budget 2021 (CD agg) (thousands)	10.4685	4.6632	3.6329	7.3987	9.2408	12.2010	32.0144	1270
FBC taxes budget 2024 (CD agg) (thousands)	15.5926	6.3300	6.6344	11.2212	13.4613	18.7518	46.1225	1270
Panel C: Wage Levels								
Annual total wages, 2021 (billions USD, CD agg.)	21.0853	16.2093	0.0005	13.1335	17.5283	23.6584	159.4335	1270
Annual total wages, 2024 (billions USD, CD agg.)	25.0538	18.6609	0.0006	15.9010	20.8778	28.1266	193.0775	1270
Annual average pay, 2021 (thousands USD, CD agg.)	62.6851	18.8076	39.5025	50.8971	58.4239	69.6285	185.2530	1270
Annual average pay, 2024 (thousands USD, CD agg.)	69.9289	19.1829	44.5637	57.7540	66.0737	77.3402	197.0430	1270
Panel D: Real Wage Growth (Average Pay – Inflation)								
Real wage growth: total (Avg Pay)	-14.9213	8.1958	-47.2128	-20.2612	-14.2713	-8.6854	3.4041	1270
Real wage growth: food (Avg Pay)	-16.2614	5.8025	-49.5415	-19.6126	-15.4161	-12.2965	0.0485	1270
Real wage growth: transportation (Avg Pay)	-15.4493	10.3629	-135.8562	-17.7780	-14.4625	-11.4876	2.6414	1270
Real wage growth: housing (Avg Pay)	-17.0362	9.0327	-47.0286	-23.0633	-15.9505	-10.9339	10.9392	1270
Real wage growth: childcare (Avg Pay)	-15.2028	32.6458	-172.9601	-30.6757	-9.7537	7.9650	40.1969	1270
Real wage growth: healthcare (Avg Pay)	-5.2749	11.3701	-74.2341	-12.1787	-5.2847	1.5471	25.6893	1270
Real wage growth: othernecessities (Avg Pay)	-7.5245	5.5959	-25.8555	-11.1614	-7.0062	-4.2915	12.0886	1270
Real wage growth: taxes (Avg Pay)	-40.0062	19.3726	-108.3072	-54.8937	-40.5666	-23.5581	-4.1502	1270

Notes: Observations are congressional district by congressional cycle in the CD estimation sample ($N \approx 1{,}270$ district-cycle observations across the 116th, 117th, and 118th Congresses). Panel A reports 2020–2024 changes in district-level U.S. House electoral outcomes computed from MIT Election Lab House returns. Panel B reports 2021–2024 percent changes in the Economic Policy Institute Family Budget Calculator costs aggregated from county to district by population-weighted average. Panel C reports 2021–2024 percent changes in total covered wages and average annual pay (QCEW), aggregated to the district. Panel D reports the category-specific wage–price gap.

confidence bands. This is what one would expect if pre-shock housing-market levels and pre-shock unemployment carry limited information beyond what the baseline covariates already capture. To the extent that a reader finds these additional regressors substantively informative, their inclusion does not materially alter the paper's conclusions.

Columns (3) and (4) add contemporaneous regressors that are themselves outcomes of the inflation–wage shock, and the coefficients move in the direction the bad-controls argument predicts. The residual price coefficient δ_2 attenuates from -0.121 to -0.099 and then to -0.079 as successive post-treatment regressors are included, with the state-clustered significance level falling from one percent to ten percent in column (4). The real-wage-growth coefficient δ_1 attenuates more sharply and becomes statistically indistinguishable from zero once the ERS unemployment rate is added. Both patterns are mechanical consequences of conditioning on variables downstream of the treatment: unemployment in 2022–2023 is, as a matter of macroeconomic accounting, closely tied to the same labor-market dynamics that determine the gap at the county level, so it partials out a large share of the gap's relevance for voting; and the FHFA House Price Index is by basket-weight argument the largest component of EPI inflation, so its inclusion removes the housing contribution to the price signal and leaves $\tilde{\delta}_2$ measuring only the non-housing components. The sign of δ_2 is preserved across all four columns under both inference procedures.

Table D.2 shows a different pattern for turnout: both δ_1 and δ_2 are positive in every specification and their point estimates rise as the additional controls are added, from 0.099 and 0.070 in the baseline to 0.202 and 0.151 in column (4). The channels that determine turnout changes between presidential elections—residential mobility, registration activity, community stability—operate alongside rather than through the wage–price gap, and are correlated with local housing dynamics and unemployment for reasons unrelated to the inflation episode itself. For the turnout regression the additional regressors therefore resolve orthogonal cross-county variation rather than partial out the mechanism of interest, and the coefficients sharpen. The sign of both coefficients is preserved in every column.

Table D.1
Wage–price gap re-parameterization: robustness to housing-price and unemployment controls — Republican two-party vote-share change, 2020–2024

	(1) Baseline	(2) + ZHVI$_{2020}$ + LAUS$_{2020}$	(3) + ERS urate 2022–2023	(4) + All additional controls
Real Wage Growth	−0.0517	−0.0614	−0.0325	−0.0203
	(0.0208)**	(0.0262)**	(0.0205)	(0.0352)
	[0.0311]*	[0.0357]*	[0.0288]	[0.0432]
Inflation	−0.1205	−0.1123	−0.0987	−0.0789
	(0.0275)***	(0.0279)***	(0.0261)***	(0.0312)**
	[0.0386]***	[0.0401]***	[0.0354]***	[0.0451]*
Observations	3,102	2,984	3,102	2,523
R^2	0.413	0.446	0.429	0.472
State FE	Yes	Yes	Yes	Yes
Baseline controls	Yes	Yes	Yes	Yes
+ ZHVI + LAUS (2020 levels)	No	Yes	No	Yes
+ ERS unemployment (2022, 2023)	No	No	Yes	Yes
+ FHFA HPI (annual change, 2021–2024)	No	No	No	Yes

Notes: Ordinary least squares estimates of the wage–price gap reparameterization of Equation (3) with total Family Budget Calculator inflation as the price measure. Observations are U.S. counties. The dependent variable is the change in the Republican two-party presidential vote share between 2020 and 2024. "Real Wage Growth" is the coefficient δ_1 on the standardized wage–price gap $G_c^{\text{total}} = \Delta W_c - \Delta P_c^{\text{total}}$; "Inflation" is the coefficient δ_2 on standardized total residual inflation. Column (1) is the paper's preferred specification with state fixed effects and the baseline covariate set. Column (2) additionally includes pre-period (December 2020) levels of the Zillow Home Value Index and the BLS Local Area Unemployment Statistics (LAUS) unemployment rate. Column (3) additionally includes the USDA Economic Research Service (ERS) unemployment rate for 2022 and 2023, matching the control set used in the main joint specification of Flavin (2026). Column (4) includes the additional regressors of columns (2) and (3) together with four annual changes of the FHFA House Price Index for 2021–2024. All continuous variables are standardized to zero mean and unit variance. Heteroskedasticity-robust standard errors in parentheses; state-clustered block-bootstrap standard errors (100 replications, seed 1) in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In either outcome, the sign of the association between local inflation, the wage–price gap, and the electoral response is preserved across all four specifications, including those that are not econometrically preferred. Where the additional regressors are proper pre-treatment covariates, as in column (2), the preferred estimates are statistically unchanged. Where the additional regressors are post-treatment outcomes of the shock under study, as in columns (3) and (4), the coefficients attenuate in the direction and magnitude that the bad-controls argument predicts, but the substantive conclusions of the paper are preserved. The specification reported in the main text remains the preferred one on econometric grounds.

Sample sizes fall to 2,984 in column (2) and 2,523 in column (4) because Zillow and FHFA county coverage is incomplete for a number of small-population counties. Columns (1) and (3) preserve the full baseline sample of 3,102 counties. Restricting column (1) to the smaller sample used in column (4) does not alter the qualitative conclusions.

E. Robustness to Partisan-Realignment Trends

A separate concern is that the negative residual-inflation coefficient is loading on a cross-county partisan trend correlated with income or urbanization. In the 2020–2024 period the U.S. realignment continued along familiar lines: lower-income and more rural counties shifted further Republican while higher-income and more urbanized counties did not, and that trend is plausibly itself correlated with the cross-section of cost-of-living growth that the EPI Family Budget Calculator records. State fixed effects do not absorb the trend if it operates within state along the income or urbanization dimension, and the resulting δ_2 might capture the within-state continuation of the partisan trend rather than a behavioral response to local prices. This appendix conditions directly on the cross-county partisan trend by adding lagged differences of the Republican two-party presidential vote share at three horizons, and asks how much of the residual-inflation coefficient survives.

Tables E.1 and E.2 report the short-horizon test. Column (1) is the paper's preferred specification. Column (2) adds the 2016–2020 county-level change in the two-party Republican presidential vote share as a direct control for the pre-existing partisan trend. The 2016 and 2020 vote-share data are drawn from the MIT Election Lab county presidential returns and rendered as the change in the two-party Republican share between the two elections, in percentage points, mirroring the construction of the dependent variable on the left-hand side.

Table D.2
Wage–price gap re-parameterization: robustness to housing-price and unemployment controls — turnout change, 2020–2024

	(1) Baseline	(2) + ZHVI_{2020} + LAUS_{2020}	(3) + ERS urate 2022–2023	(4) + All additional controls
Real Wage Growth	0.0987	0.1329	0.1083	0.2020
	(0.0489)**	(0.0333)***	(0.0514)**	(0.0406)***
	[0.0618]	[0.0422]***	[0.0616]*	[0.0512]***
Inflation	0.0703	0.0979	0.0780	0.1511
	(0.0406)*	(0.0292)***	(0.0422)*	(0.0328)***
	[0.0516]	[0.0437]**	[0.0505]	[0.0443]***
Observations	3,102	2,984	3,102	2,523
R^2	0.402	0.425	0.404	0.465
State FE	Yes	Yes	Yes	Yes
Baseline controls	Yes	Yes	Yes	Yes
+ ZHVI + LAUS (2020 levels)	No	Yes	No	Yes
+ ERS unemployment (2022, 2023)	No	No	Yes	Yes
+ FHFA HPI (annual change, 2021–2024)	No	No	No	Yes

Notes: Ordinary least squares estimates of the wage–price gap reparameterization of Equation (3) with total Family Budget Calculator inflation as the price measure. Observations are U.S. counties. The dependent variable is the change in county-level presidential turnout between 2020 and 2024. "Real Wage Growth" is the coefficient δ_1 on the standardized wage–price gap; "Inflation" is the coefficient δ_2 on standardized total residual inflation. Column structure matches Table D.1: Column (1) is the preferred specification with state fixed effects and the baseline covariate set; Column (2) adds pre-period (December 2020) Zillow Home Value Index and LAUS unemployment-rate levels; Column (3) adds the USDA ERS unemployment rate for 2022 and 2023 used by Flavin (2026); Column (4) stacks these augmentations together with four annual changes of the FHFA House Price Index for 2021–2024. All continuous variables are standardized to zero mean and unit variance. Heteroskedasticity-robust standard errors in parentheses; state-clustered block-bootstrap standard errors (100 replications, seed 1) in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The 2016–2020 pre-trend itself enters with a large positive coefficient on the Republican two-party share change (0.302 in column (2) of Table E.1, $p < 0.01$ under both inference procedures): a county that shifted Republican between 2016 and 2020 by one standard deviation continues to shift Republican between 2020 and 2024 by roughly 0.30 of a standard deviation. The realignment has substantial period-over-period momentum. The pre-trend control, however, does not subsume the residual-inflation coefficient: δ_2 moves from -0.121 in the baseline to -0.117, a three-percent attenuation that is well within the confidence band of the baseline. The 2016–2020 partisan-trend hypothesis, in this direct form, does not explain the cross-county inflation–vote relationship. The wage–price gap coefficient δ_1 is attenuated to -0.047 and remains significant at the ten-percent level under heteroskedasticity-robust inference, though it falls below the conventional threshold under state-clustered inference.

A natural concern is that the four-year window does not capture the full diploma realignment, which has been visible since at least 2008 and accelerated through Trump's first term. Tables E.3 and E.4 address this with two longer-horizon specifications. Column (2) adds the 2008–2020 cross-county Republican two-party share change as a single twelve-year secular-trend control. Column (3) decomposes the long pre-trend into its 2008–2016 (Obama-era) and 2016–2020 (late-Trump) segments and includes both.

The long-horizon pre-trend is large and highly significant. The standardized 2008–2020 Republican-share change enters column (2) of Table E.3 with a coefficient of 0.421 ($p < 0.01$ under both inference procedures); each one-standard-deviation Republican shift over 2008–2020 predicts a 0.42-standard-deviation continuation over 2020–2024. Decomposed in column (3), the Obama-era segment carries 0.216 and the late-Trump segment carries 0.289, both at the one-percent level. The realignment is real, secular, and quantitatively important.

The diagnostic finding, however, is that this long-horizon partisan trend does not absorb the residual-inflation coefficient. In column (2) of Table E.3, δ_2 moves from -0.121 to -0.117, a three-percent attenuation that is well within the confidence band of the baseline. The decomposition in column (3) gives essentially the same number: $\delta_2 = -0.117$. Even with both segments of the 12-year secular realignment in the regression, the negative inflation coefficient is unchanged. The wage–price gap coefficient δ_1 is also preserved: -0.048 in column (2) and -0.047 in column (3), both significant at the ten-percent level under both inference procedures.

This is a strong negative result for the partisan-realignment hypothesis as currently characterized. The 12-year cross-county Republican trend does not subsume the 2020–2024 cross-county inflation–vote relationship. Whatever drives the negative sign of δ_2 is not a continuation of the secular diploma realignment in any horizon we can capture

Table E.1
Wage–price gap re-parameterization: short-horizon partisan-trend robustness — Republican two-party vote-share change, 2020–2024

	(1) Baseline	(2) + 2016–2020 pre-trend
Real Wage Growth	−0.0517	−0.0468
	(0.0208)**	(0.0270)*
	[0.0311]*	[0.0295]
Inflation	−0.1205	−0.1169
	(0.0275)***	(0.0296)***
	[0.0386]***	[0.0349]***
Pre-trend (2016–2020)		0.3021
		(0.0297)***
		[0.0436]***
Observations	3,102	3,102
R^2	0.413	0.457
Baseline controls	Yes	Yes
State FE	Yes	Yes
2016–2020 pre-trend	No	Yes

Notes: Ordinary least squares estimates of the wage–price gap reparameterization of Equation (3) with total Family Budget Calculator inflation as the price measure. Observations are U.S. counties. The dependent variable is the change in the Republican two-party presidential vote share between 2020 and 2024. "Real Wage Growth" is the coefficient δ_1 on the standardized wage–price gap G_c^{total}; "Inflation" is the coefficient δ_2 on standardized total residual inflation. "Pre-trend (2016–2020)" is the coefficient on the standardized 2016–2020 county-level change in the two-party Republican presidential vote share, constructed from the MIT Election Lab county presidential returns. Column (1) is the paper's preferred specification with state fixed effects and the baseline covariate set. Column (2) adds the 2016–2020 Republican two-party vote-share change as a partisan-trend control. All continuous variables are standardized to zero mean and unit variance. Heteroskedasticity-robust standard errors in parentheses; state-clustered block-bootstrap standard errors (100 replications, seed 1) in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

with vote-share lagged differences. It is a 2020–2024-specific within-state cross-county phenomenon. Consistent with the rest of the paper, we interpret the unexplained component as residual demand-pressure salience: counties where the 2021–2024 nominal economy was running hotter, conditional on the wage–price gap, voted relatively more Democratic than the 12-year secular trend would have predicted. Section 7.1 develops the interpretation; here we simply note that the within-state cross-county object is robust to partisan-trend conditioning at every horizon.

For the turnout outcome (Tables E.2 and E.4), the residual-inflation coefficient is positive and weakly significant in every specification. The 2016–2020 pre-trend is negatively associated with the turnout change (-0.144, $p < 0.01$): counties that swung Republican between 2016 and 2020 had *lower* turnout growth between 2020 and 2024. The Obama-era segment of the long pre-trend, by contrast, is positive and small. Counties that shifted Republican specifically over the late-Trump period had lower turnout growth in 2024, while pre-2016 partisan momentum did not.

F. Geographic Distribution Maps

Figures F.1 to F.5 present the geographical distribution of the main variables in the paper.

G. Testing for Non-Linearities in the Inflation–Voting Relationship

A central finding of the main analysis is that local inflation retains a negative association with Republican vote gains even after conditioning on the wage–price gap (the coefficient δ_2 in equation 3). This negative residual price effect, documented in Table 3 and discussed in Section 6, is difficult to reconcile with a strict retrospective pocketbook model in which voters respond only to changes in real wages.

One possible explanation is that the linear specification imposes too restrictive a functional form. If the political response to inflation differs at different levels of price growth (for example, moderate inflation in economically dynamic areas versus extreme inflation that triggers punishment voting), the linear coefficient averages over these heterogeneous effects. Composition effects could also matter: if the negative linear slope is driven by the dominance of low-inflation rural counties where Republicans gained the most, the relationship might flatten or reverse at higher inflation levels.

Table E.2
Wage–price gap re-parameterization: short-horizon partisan-trend robustness — turnout change, 2020–2024

	(1) Baseline	(2) + 2016–2020 pre-trend
Real Wage Growth	0.0987	0.0964
	(0.0489)**	(0.0452)**
	[0.0618]	[0.0609]
Inflation	0.0703	0.0686
	(0.0406)*	(0.0379)*
	[0.0516]	[0.0507]
Pre-trend (2016–2020)		−0.1436
		(0.0285)***
		[0.0255]***
Observations	3,102	3,102
R^2	0.402	0.412
Baseline controls	Yes	Yes
State FE	Yes	Yes
2016–2020 pre-trend	No	Yes

Notes: Ordinary least squares estimates of the wage–price gap reparameterization. Observations are U.S. counties. The dependent variable is the change in turnout between 2020 and 2024. "Real Wage Growth" is the coefficient δ_1 on the standardized wage–price gap; "Inflation" is the coefficient δ_2 on standardized total residual inflation. "Pre-trend (2016–2020)" is the coefficient on the standardized 2016–2020 county-level change in the two-party Republican presidential vote share. Column (1) is the paper's preferred specification. Column (2) adds the 2016–2020 partisan-trend control. All continuous variables are standardized to zero mean and unit variance. Heteroskedasticity-robust standard errors in parentheses; state-clustered block-bootstrap standard errors (100 replications, seed 1) in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

This appendix tests whether allowing for flexible functional forms alters the interpretation. We employ three complementary approaches: residualized binscatter plots, natural cubic spline regressions with formal F-tests for linearity, and heatmaps of the joint inflation–wage distribution. All specifications condition on state fixed effects and the full set of county-level controls used in the main analysis. The estimation sample comprises 3,102 counties.

G.1. Residualized binscatter evidence

We begin with residualized binscatter plots, which provide a non-parametric visualization of the conditional expectation of each electoral outcome given inflation or the wage–price gap, after partialing out state fixed effects and controls. Each plot groups the residualized explanatory variable into 20 quantile bins, computes the mean of the residualized outcome within each bin, and overlays both a linear fit (dashed gray) and a quadratic fit with confidence band (orange). The linear slope and its standard error are reported on each panel.

For each outcome, three variants are shown:

1. **Baseline inflation.** Outcome versus inflation, both residualized with respect to state fixed effects and the full controls. This visualizes the relationship underlying β_1 in the baseline specification (1).

2. **Wage–price gap.** Outcome versus $G_i^{(k)}$, using the same controls. This visualizes the real wage growth channel (δ_1).

3. **Inflation conditional on the gap (δ_2).** Outcome versus inflation, residualized with respect to controls *and* the wage–price gap. This is the key diagnostic: it shows whether the residual inflation–voting relationship exhibits non-linear features that could alter the interpretation of the negative δ_2.

Across all outcomes and specifications, the binscatter plots show that the conditional expectation of electoral outcomes given inflation is largely monotone and approximately linear. The quadratic fits track the linear fits closely, without systematic curvature that would overturn the linear interpretation. This is particularly true for vote share and vote margin, where the bin means align tightly with the linear prediction. For turnout, there is slightly more dispersion in the tails, but no clear non-monotonic pattern. The δ_2 panels, which display inflation after partialing out the wage–price gap, show the same approximately linear, negative relationship as the baseline panels, indicating that the negative

Table E.3
Wage–price gap re-parameterization: long-horizon partisan-trend robustness — Republican two-party vote-share change, 2020–2024

	(1) Baseline	(2) + 2008–2020 pre-trend	(3) Stacked lagged differences
Real Wage Growth	−0.0517	−0.0476	−0.0465
	(0.0208)**	(0.0198)**	(0.0242)*
	[0.0311]*	[0.0262]*	[0.0267]*
Inflation	−0.1205	−0.1174	−0.1166
	(0.0275)***	(0.0264)***	(0.0281)***
	[0.0386]***	[0.0356]***	[0.0342]***
Pre-trend (2008–2020)		0.4206	
		(0.0298)***	
		[0.0491]***	
Pre-trend (2008–2016)			0.2162
			(0.0361)***
			[0.0499]***
Pre-trend (2016–2020)			0.2888
			(0.0305)***
			[0.0405]***
Observations	3,102	3,101	3,101
R^2	0.413	0.454	0.467
Baseline controls	Yes	Yes	Yes
State FE	Yes	Yes	Yes
2008–2020 pre-trend	No	Yes	No
2008–2016 + 2016–2020	No	No	Yes

Notes: Ordinary least squares estimates of the wage–price gap reparameterization of Equation (3) with total Family Budget Calculator inflation as the price measure. Observations are U.S. counties. The dependent variable is the change in the Republican two-party presidential vote share between 2020 and 2024. The pre-trend variables are standardized changes in the county-level Republican two-party presidential vote share over the indicated horizons, constructed from the MIT Election Lab county presidential returns. Column (1) is the paper's preferred specification with state fixed effects and the baseline covariate set. Column (2) adds the 2008–2020 long-horizon pre-trend. Column (3) decomposes that pre-trend into the 2008–2016 (Obama-era) and 2016–2020 (late-Trump) segments and includes both. All continuous variables are standardized to zero mean and unit variance. Heteroskedasticity-robust standard errors in parentheses; state-clustered block-bootstrap standard errors (100 replications, seed 1) in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

association documented in the main tables reflects a consistent downward slope across the inflation distribution rather than an artifact of averaging over a non-linear conditional expectation.

Table E.4
Wage–price gap re-parameterization: long-horizon partisan-trend robustness — turnout change, 2020–2024

	(1) Baseline	(2) + 2008–2020 pre-trend	(3) Stacked lagged differences
Real Wage Growth	0.0987	0.0949	0.0939
	(0.0489)**	(0.0482)**	(0.0437)**
	[0.0618]	[0.0607]	[0.0588]
Inflation	0.0703	0.0676	0.0669
	(0.0406)*	(0.0401)*	(0.0368)*
	[0.0516]	[0.0512]	[0.0501]
Pre-trend (2008–2020)		−0.0732	
		(0.0304)**	
		[0.0404]*	
Pre-trend (2008–2016)			0.0667
			(0.0310)**
			[0.0635]
Pre-trend (2016–2020)			−0.1476
			(0.0286)***
			[0.0252]***
Observations	3,102	3,101	3,101
R^2	0.402	0.403	0.412
Baseline controls	Yes	Yes	Yes
State FE	Yes	Yes	Yes
2008–2020 pre-trend	No	Yes	No
2008–2016 + 2016–2020	No	No	Yes

Notes: Ordinary least squares estimates of the wage–price gap reparameterization. Observations are U.S. counties. The dependent variable is the change in turnout between 2020 and 2024. The pre-trend variables are standardized changes in the county-level Republican two-party presidential vote share over the indicated horizons. Column (1) is the paper's preferred specification. Column (2) adds the 2008–2020 long-horizon pre-trend. Column (3) decomposes that pre-trend into the 2008–2016 and 2016–2020 segments. All continuous variables are standardized to zero mean and unit variance. Heteroskedasticity-robust standard errors in parentheses; state-clustered block-bootstrap standard errors (100 replications, seed 1) in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure F.1: County-level electoral outcomes — changes, 2020–2024

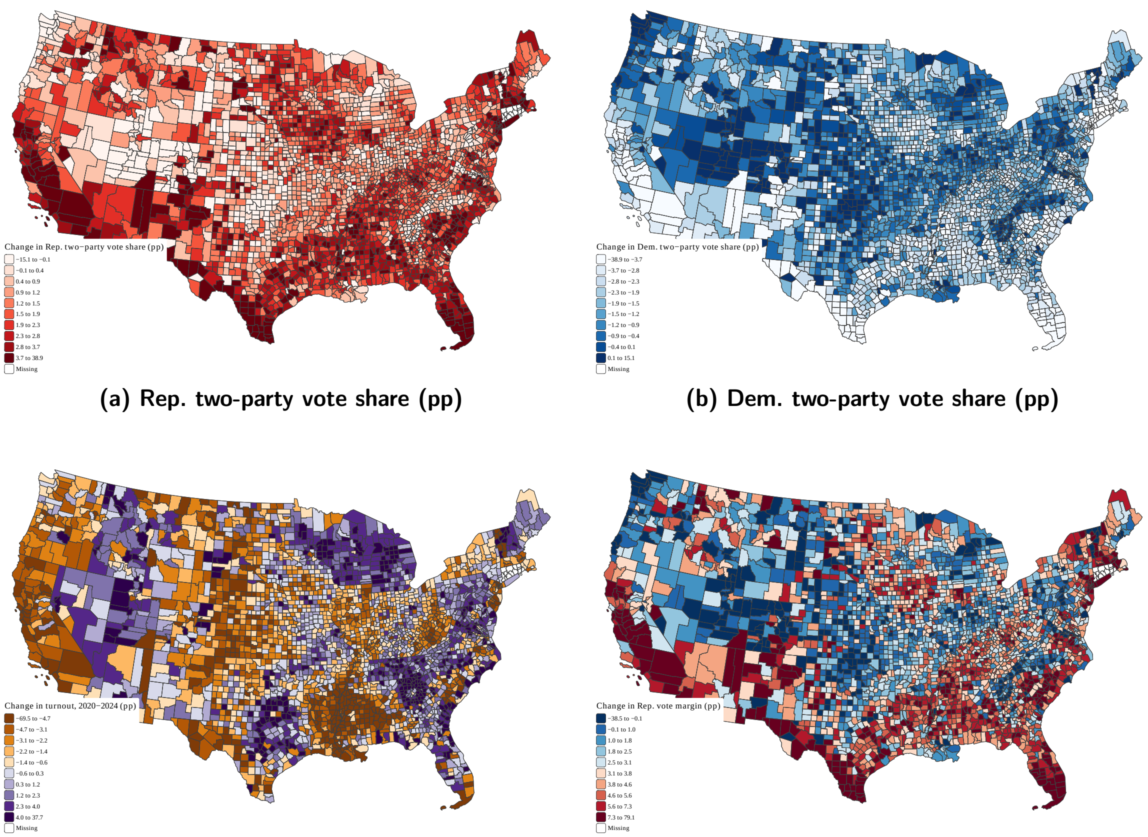

(a) Rep. two-party vote share (pp) **(b) Dem. two-party vote share (pp)**

(c) Turnout (pp) **(d) Rep. vote margin (pp)**

Notes: Each map shows the county-level change in the indicated electoral outcome between the 2020 and 2024 presidential elections, corresponding to Panel A of Table 1. Counties are shaded by decile on a diverging scale. CONUS counties only (Alaska, Hawaii, and territories excluded). Albers Equal-Area projection (EPSG:5070).

Figure F.2: Total wage growth and total inflation, 2021–2024

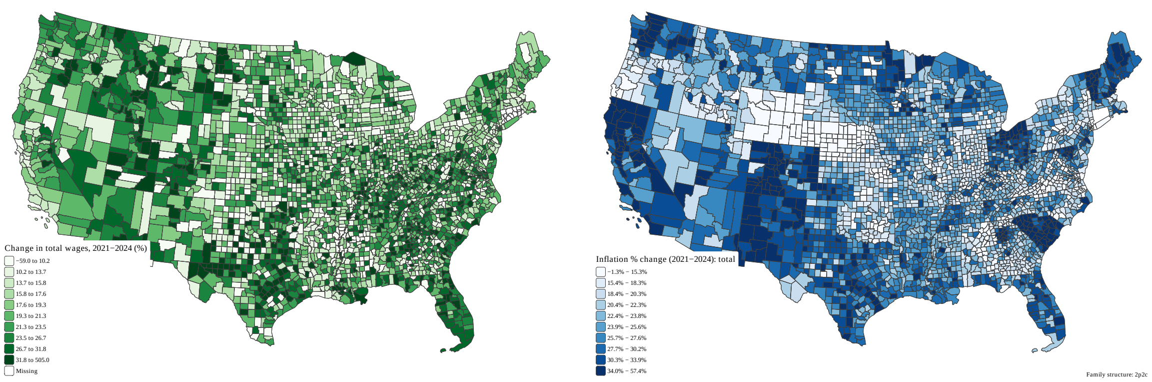

(a) Change in wages (%) **(b) Total inflation (%)**

Notes: Panel (a) shows the percent change in total covered wages (QCEW) between 2021 and 2024, corresponding to Panel C of Table 1. Panel (b) shows the percent change in total family budget costs (EPI FBC, two-parent two-child family) over the same period (same as Figure 2(b)). Counties are shaded by decile. CONUS counties only. Albers Equal-Area projection (EPSG:5070).

Figure F.3: County-level electoral outcomes — levels

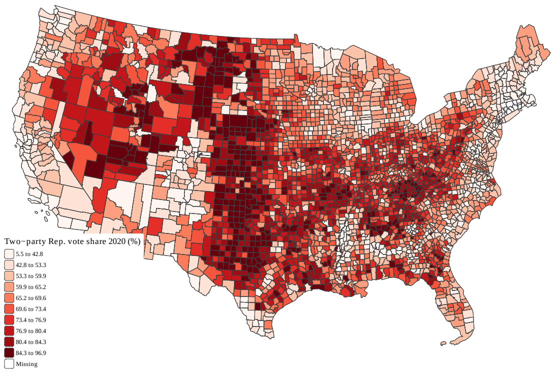

(a) Two-party Rep. vote share, 2020

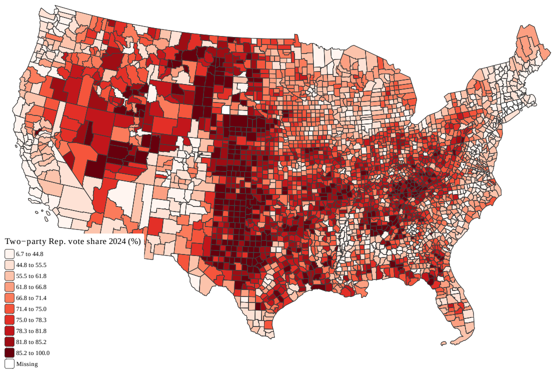

(b) Two-party Rep. vote share, 2024

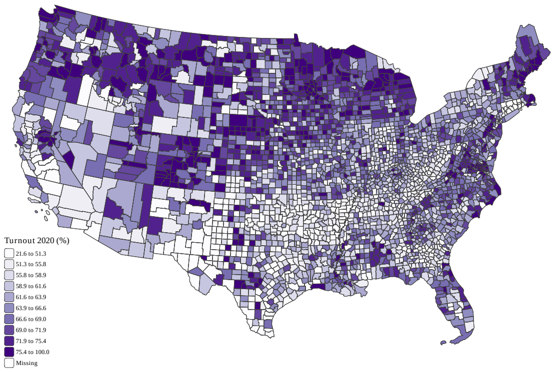

(c) Turnout, 2020

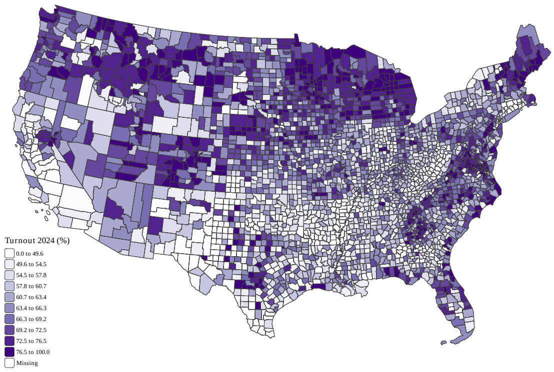

(d) Turnout, 2024

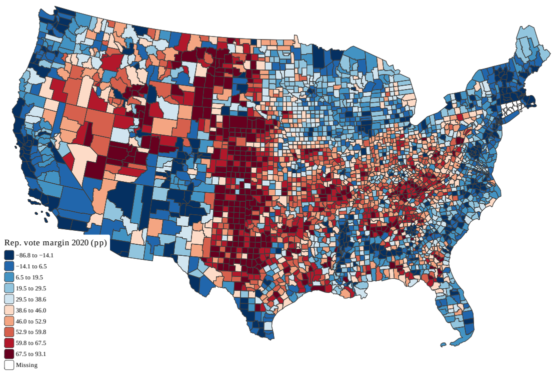

(e) Rep. vote margin, 2020

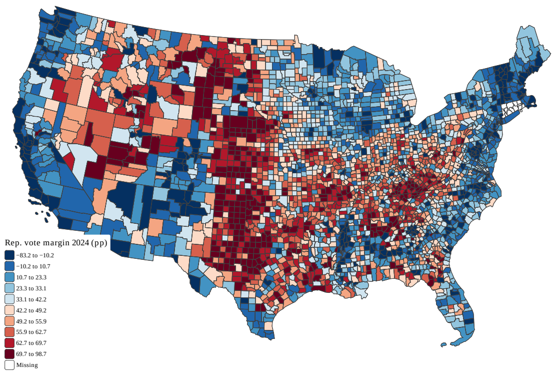

(f) Rep. vote margin, 2024

Notes: Each map shows the county-level distribution of the indicated electoral outcome variable, corresponding to Panel A of Table C.3. Counties are shaded by decile. CONUS counties only (Alaska, Hawaii, and territories excluded). Albers Equal-Area projection (EPSG:5070).

Figure F.4: County-level inflation, 2021–2024, by expenditure category

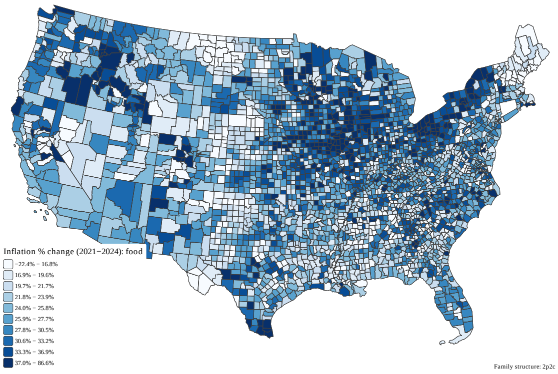

(a) Food

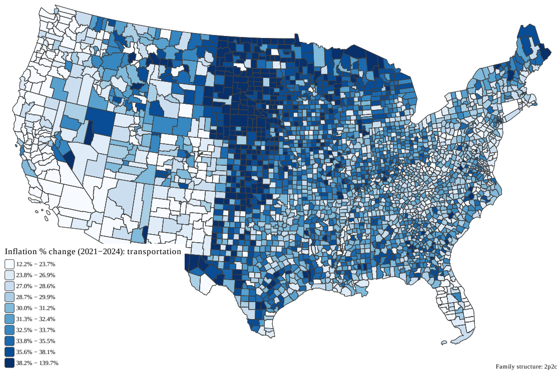

(b) Transportation

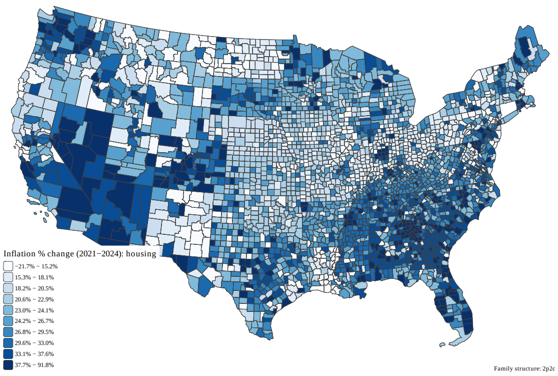

(c) Housing

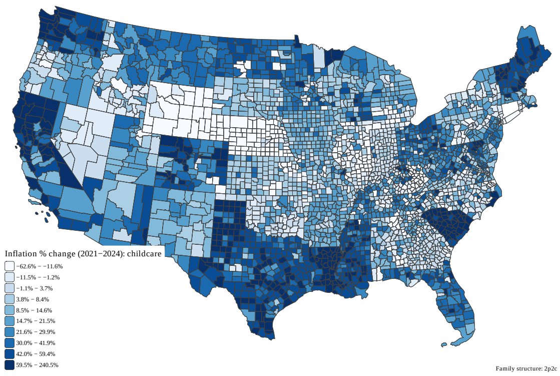

(d) Childcare

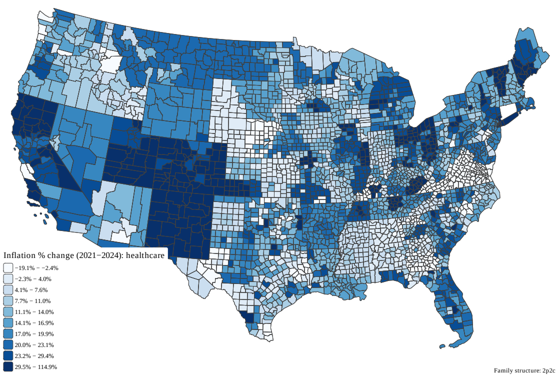

(e) Healthcare

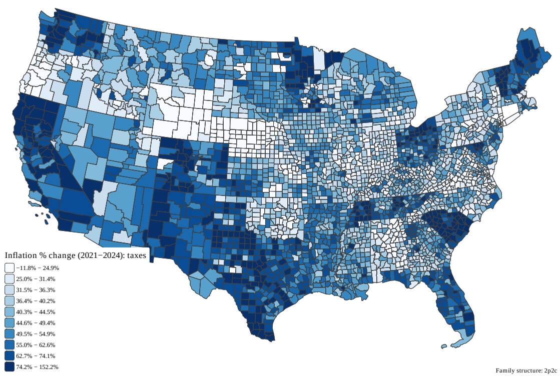

(f) Taxes

Notes: Each map shows the percent change in the EPI Family Budget Calculator cost estimate (two-parent, two-child family) between 2021 and 2024 for the indicated expenditure category, corresponding to Panel B of Table 1. Counties are shaded by decile. CONUS counties only. Albers Equal-Area projection (EPSG:5070).

Figure F.5: County-level real wage growth (total wages), by expenditure category

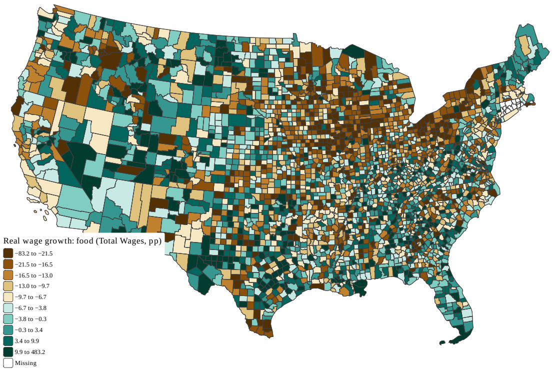

(a) Food

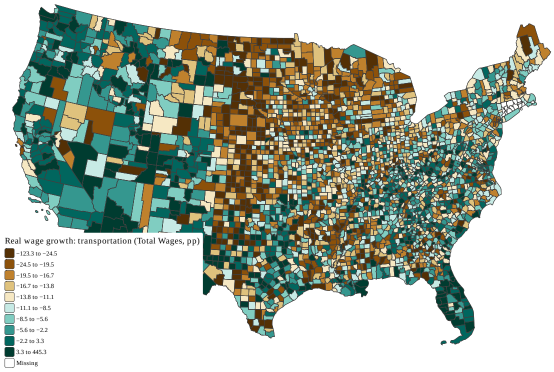

(b) Transportation

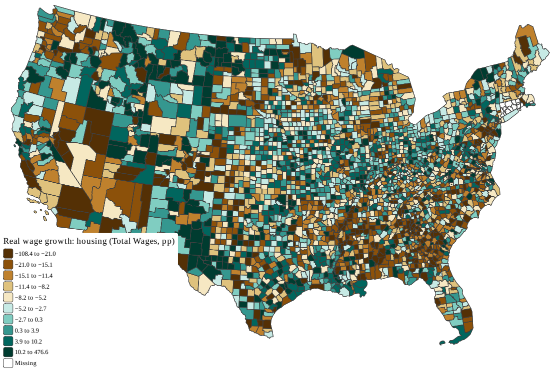

(c) Housing

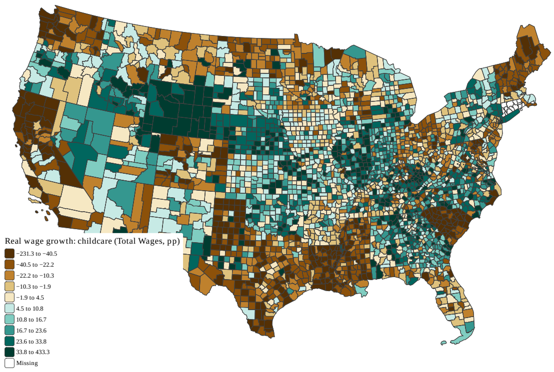

(d) Childcare

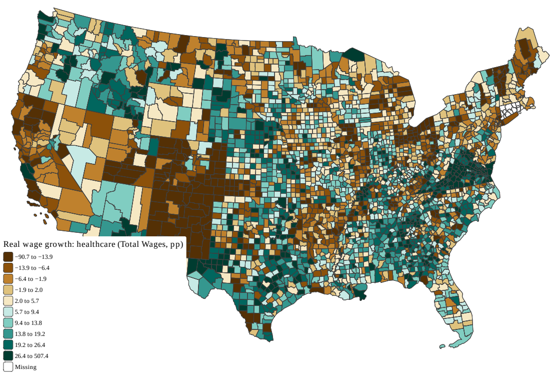

(e) Healthcare

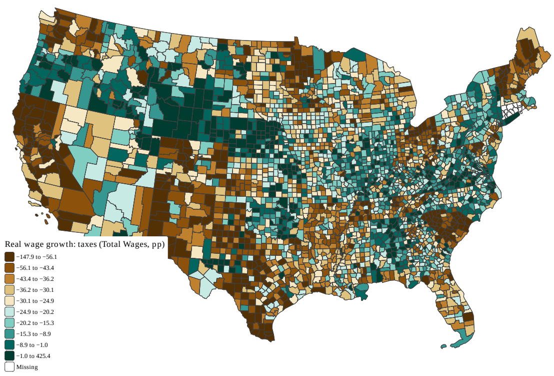

(f) Taxes

Notes: Each map shows the wage–price gap $G_i^{(k)} = \Delta W_i^{\text{total}} - \Delta P_i^{(k)}$ for the indicated expenditure category, corresponding to Panel D of Table 1. Total wages are from the QCEW. Counties are shaded by decile on a diverging scale. CONUS counties only. Albers Equal-Area projection (EPSG:5070).

Figure G.1: Vote share change vs. total inflation

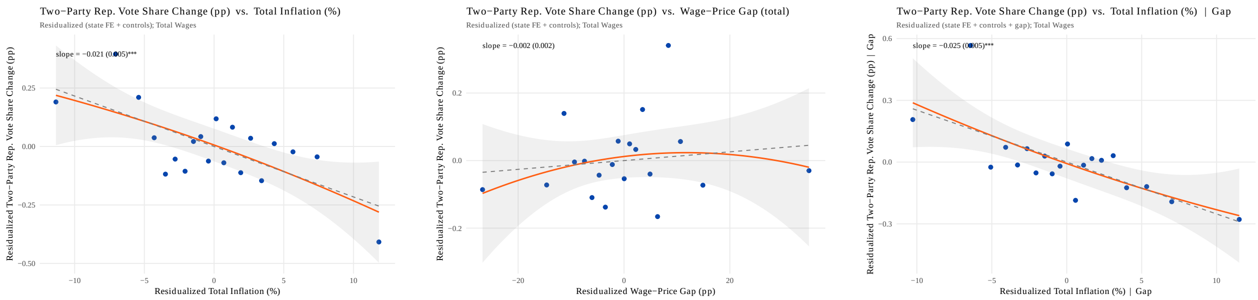

(a) Baseline (b) Gap (c) Inflation | Gap

Notes: 20 quantile bins. Gray dashed line: linear fit. Orange curve: quadratic fit with 95% confidence band. All variables residualized against state fixed effects and the full control set. Panel (c) additionally partials out the wage–price gap.

Figure G.2: Vote share change vs. food inflation

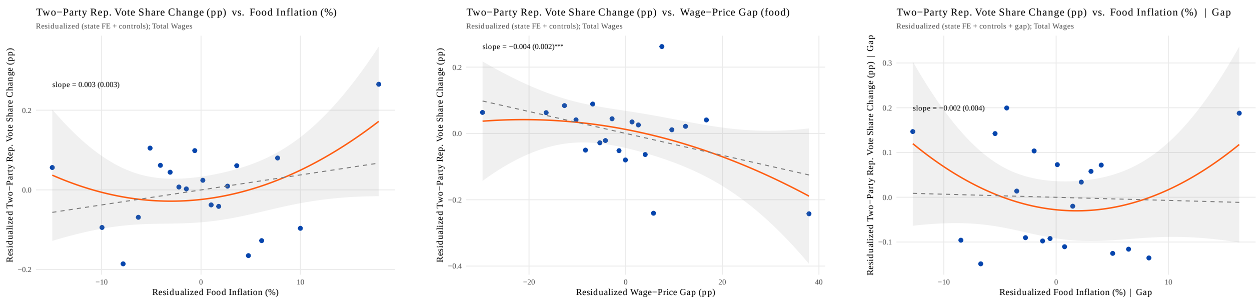

(a) Baseline (b) Gap (c) Inflation | Gap

Notes: Same as Figure G.1 but using food inflation.

Figure G.3: Vote margin change vs. total inflation

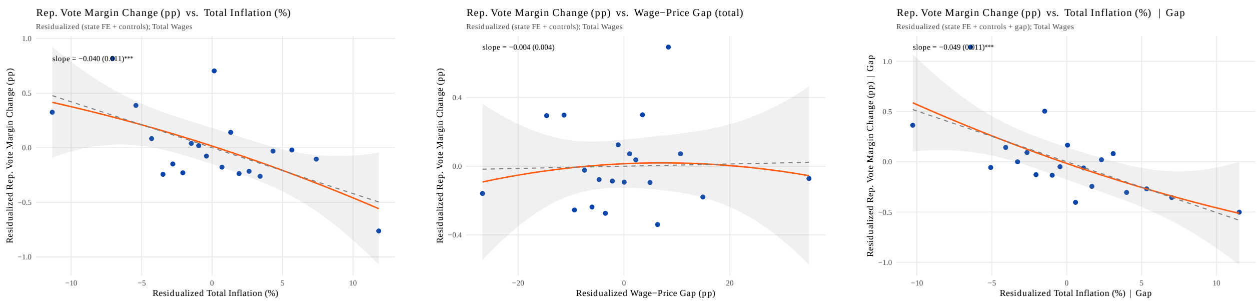

(a) Baseline (b) Gap (c) Inflation | Gap

Notes: Same as Figure G.1 but for Republican vote margin change.

Figure G.4: Vote margin change vs. food inflation

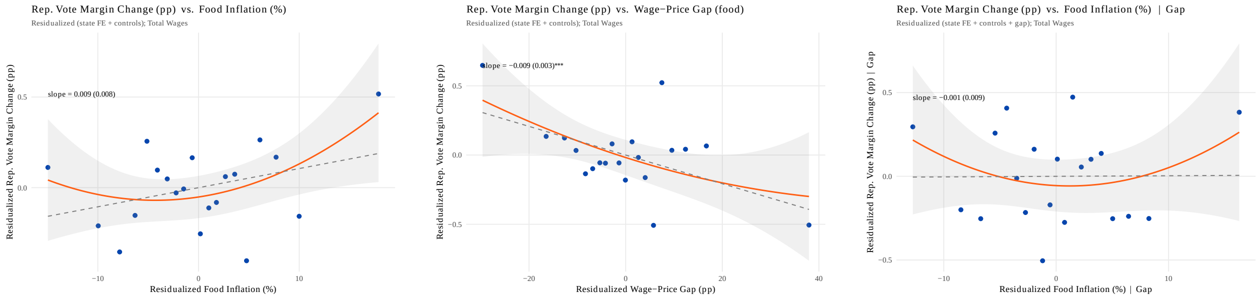

(a) Baseline (b) Gap (c) Inflation | Gap

Notes: Same as Figure G.1 but for vote margin change and food inflation.

Figure G.5: Turnout change vs. total inflation

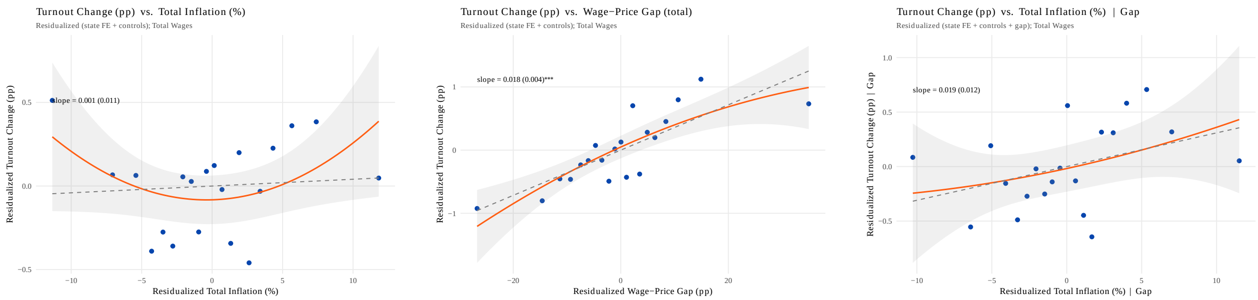

(a) Baseline (b) Gap (c) Inflation | Gap

Notes: Same as Figure G.1 but for turnout change.

Figure G.6: Turnout change vs. food inflation

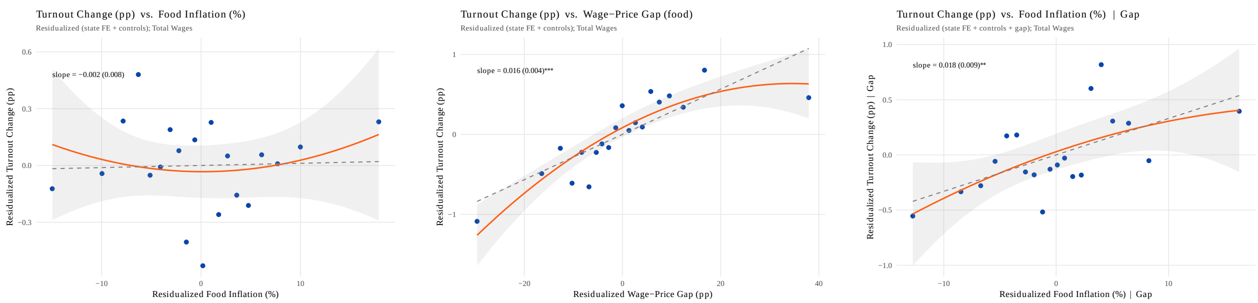

(a) Baseline (b) Gap (c) Inflation | Gap

Notes: Same as Figure G.1 but for turnout change and food inflation.

G.2. Spline regressions and formal tests of linearity

To complement the visual evidence, we estimate a natural cubic spline specification that allows the inflation–voting relationship to take a flexible, smooth non-linear form. Specifically, we replace the linear inflation term in specification (3) with a natural cubic spline basis with four degrees of freedom:

$$\Delta y_i = f_{\text{spline}}\left(\Delta P_i^{\text{total}}\right) + \delta_1 G_i^{\text{total}} + \mathbf{X}_i'\gamma + \mu_{s(i)} + \varepsilon_i,$$

where $f_{\text{spline}}(\cdot)$ is estimated from the data. The linear specification is nested within the spline model, so we use an F-test of the null hypothesis that the non-linear basis terms are jointly zero.

Figure G.7 displays the estimated spline curves (with 95% pointwise confidence intervals) alongside the linear fit for each outcome. A rug plot along the horizontal axis shows the density of the inflation distribution.

Figure G.7: Spline marginal effects: electoral outcomes vs. total inflation

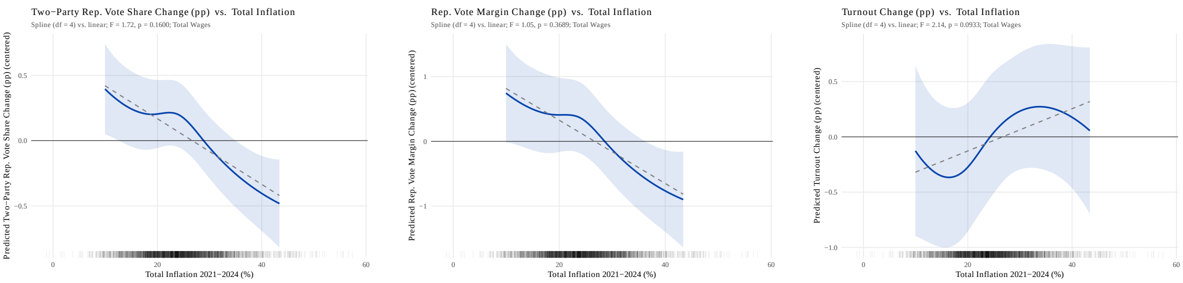

(a) Vote Share Change (b) Vote Margin Change (c) Turnout Change

Notes: Blue line: natural cubic spline (df = 4) with 95% pointwise confidence interval. Gray dashed line: linear fit. Both specifications include the wage–price gap, state fixed effects, and the full control set. Rug plot shows the marginal distribution of total inflation. Predictions are centered at zero for comparability.

Table G.1 reports the F-statistics and p-values.

The tests provide no evidence against linearity for the partisan vote outcomes: $p = 0.160$ for vote share and $p = 0.369$ for vote margin. For turnout there is marginal evidence of non-linearity ($p \approx 0.09$), consistent with slightly greater dispersion in the binscatter tails for that outcome. Visually, the spline curves in Figure G.7 are nearly indistinguishable from the linear fit across the bulk of the inflation distribution. Any deviations are modest, concentrated in the sparse tails (as indicated by the rug plot), and fall within the confidence bands. The R^2 improvement from the spline is negligible (at most 0.001).

The formal tests therefore do not reject linearity at conventional significance levels for the partisan vote outcomes that are central to the negative δ_2. For turnout, the marginal evidence does not alter the sign or substantive interpretation of the relationship.

Table G.1
F-tests of linearity: spline vs. linear specification

Outcome	F-statistic	p-value	R^2 (linear)	R^2 (spline)
Two-Party Rep. Vote Share Change	1.724	0.160	0.4126	0.4136
Rep. Vote Margin Change	1.051	0.369	0.3697	0.3703
Turnout Change	2.139	0.093	0.4021	0.4033

Notes: F-test of the null that a natural cubic spline (df = 4) in total inflation reduces to a linear term. All specifications include the wage–price gap, state fixed effects, and the full control set. $N = 3{,}102$ counties.

G.3. Joint distribution of inflation and wage growth

The final exercise examines whether the inflation–voting relationship depends on the *joint* distribution of inflation and wage growth. If non-linear interactions between prices and wages drove the negative δ_2, we would expect to see threshold effects or sign reversals in specific regions of the two-dimensional distribution, for example, Republican gains concentrated in "high inflation, low wage growth" cells.

We construct a 10×10 grid by binning counties into deciles of total inflation and total wage growth. Within each cell, we compute the mean of the residualized outcome (partialing out state fixed effects and controls). Cells are colored on a diverging scale centered at zero: blue indicates a relative shift away from Republicans, red indicates a shift toward Republicans.

Figure G.8: Inflation × wage growth heatmaps: residualized electoral outcomes

(a) Vote Share Change (b) Vote Margin (c) Turnout Change

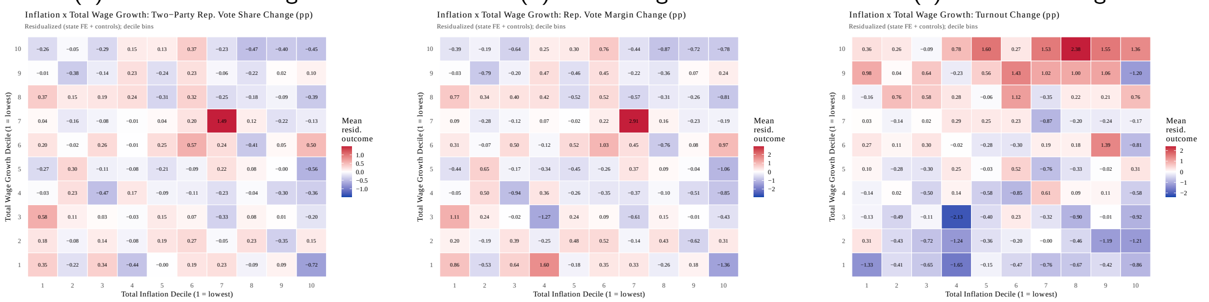

Notes: Each cell shows the mean residualized outcome (partialing out state fixed effects and the full control set) within the intersection of a total inflation decile (horizontal axis) and a total wage growth decile (vertical axis). Blue: relative shift away from Republicans. Red: relative shift toward Republicans. $N = 3{,}102$ counties.

The heatmaps do not reveal sharp thresholds or strong interaction patterns. For vote share and vote margin (panels a and b), the color gradients are generally smooth along both axes, consistent with the additive, approximately linear effects estimated in Tables 3 and 4. There is no sign reversal in the "high inflation, low wage growth" corner that a non-linear punishment-voting story would predict. For turnout (panel c), the gradient along the wage axis is more pronounced—consistent with the stronger real wage growth channel for participation documented in Table 4—but the pattern remains smooth and does not suggest discrete thresholds or strong inflation–wage interactions.

G.4. Non-linearity checks with average annual pay

Appendix H documents that the turnout–real-wage-growth association vanishes when wages are measured by average annual pay rather than total covered wages. A natural question is whether this null result is an artifact of the linear functional form. We replicate all three non-linearity exercises—binscatter plots, spline regressions, and heatmaps—using the average-pay wage–price gap in place of the total-wages gap. Because the baseline binscatter (outcome versus inflation, partialing out state fixed effects and controls only) does not involve the wage measure, those panels are mechanically identical to the total-wages results above; the gap and δ_2 panels differ.

The results are virtually identical to the total-wages case. For partisan outcomes, the F-tests do not reject linearity ($p = 0.159$ and $p = 0.364$; Table G.2), binscatter plots show monotone conditional expectations, and heatmaps

Figure G.9: Vote share change vs. total inflation (average pay)

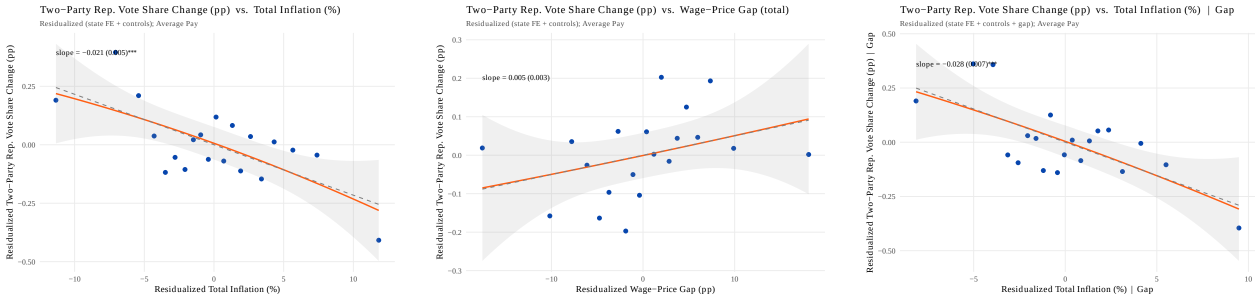

(a) Baseline (b) Gap (c) Inflation | Gap

Notes: Same as Figure G.1 but using average annual pay (QCEW) for the wage–price gap. Panel (a) is mechanically identical to the total-wages version.

Figure G.10: Vote margin change vs. total inflation (average pay)

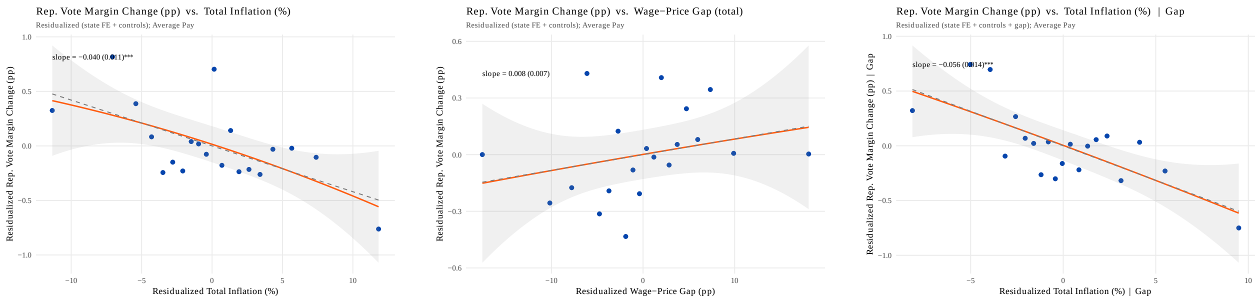

(a) Baseline (b) Gap (c) Inflation | Gap

Notes: Same as Figure G.3 but using average annual pay (QCEW) for the wage–price gap.

Figure G.11: Turnout change vs. total inflation (average pay)

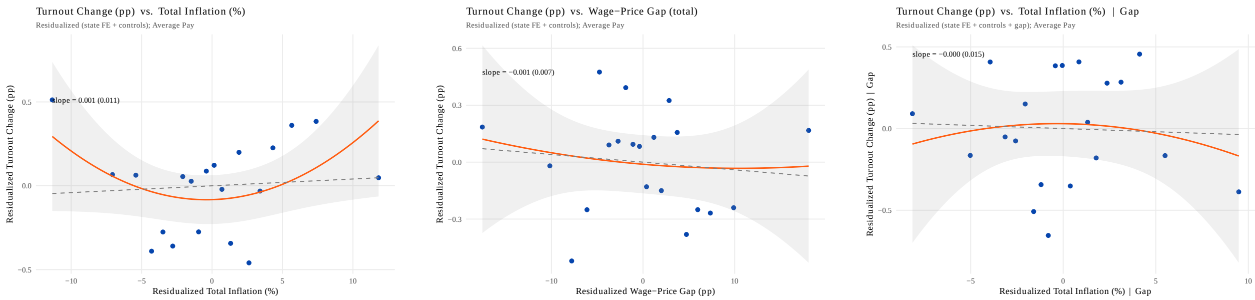

(a) Baseline (b) Gap (c) Inflation | Gap

Notes: Same as Figure G.5 but using average annual pay (QCEW) for the wage–price gap. The gap panel (b) shows a near-zero slope, confirming the null turnout–real-wage-growth association documented in Table H.2.

display smooth gradients. For turnout, the gap panel of the binscatter (Figure G.11b) confirms that the near-zero linear association documented in Table H.2 is not masking a non-linear relationship: the bin means scatter around zero with no discernible slope or curvature. The spline F-test for turnout is marginally significant ($p = 0.094$), as in the total-wages case, but the spline curve remains nearly flat, and the R^2 improvement is negligible (0.001). The heatmap (Figure G.13c) shows no gradient along the average-pay axis for turnout, in contrast to the clear wage gradient visible in the total-wages heatmap (Figure G.8c). This confirms that the null turnout result with average pay is robust to flexible functional forms: it reflects a genuine absence of association between per-worker real wage growth and turnout, not a specification artifact.

Figure G.12: Spline marginal effects: electoral outcomes vs. total inflation (average pay)

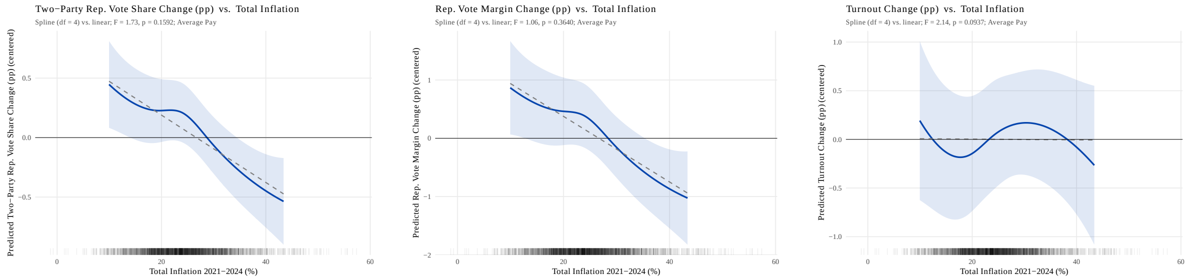

(a) Vote Share Change (b) Vote Margin Change (c) Turnout Change

Notes: Same as Figure G.7 but using average annual pay (QCEW) for the wage–price gap.

Table G.2

F-tests of linearity: spline vs. linear specification (average pay)

Outcome	F-statistic	p-value	R^2 (linear)	R^2 (spline)
Two-Party Rep. Vote Share Change	1.727	0.159	0.4119	0.4129
Rep. Vote Margin Change	1.062	0.364	0.3692	0.3698
Turnout Change	2.136	0.094	0.3974	0.3986

Notes: Same as Table G.1 but using average annual pay (QCEW) for the wage–price gap. $N = 3{,}102$ counties.

Figure G.13: Inflation × average pay growth heatmaps: residualized electoral outcomes

(a) Vote Share Change (b) Vote Margin (c) Turnout Change

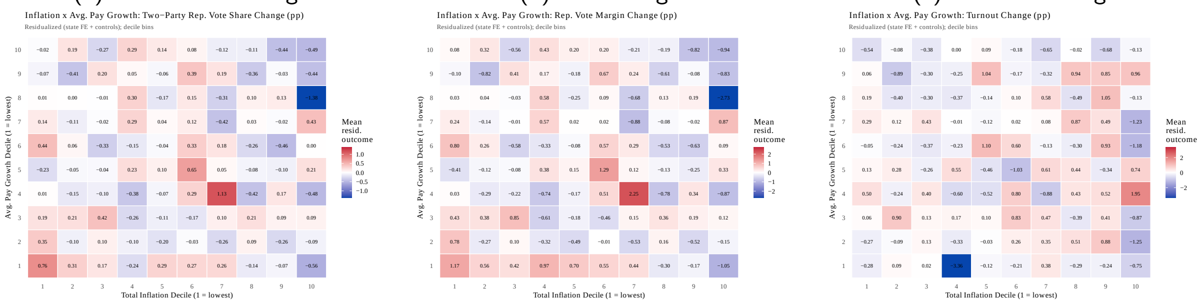

Notes: Same as Figure G.8 but using average annual pay growth (QCEW) on the vertical axis instead of total wage growth. $N = 3{,}102$ counties.

Across all exercises and both wage measures, there is little evidence that non-linearities explain the negative coefficient on inflation documented in the main results. Residualized binscatter plots show monotone, approximately linear conditional expectations across the full distribution, including in the specification that conditions on the wage–price gap. Spline regressions do not reject linearity at conventional significance levels for the partisan vote outcomes ($p = 0.160$ and $p = 0.369$ with total wages; $p = 0.159$ and $p = 0.364$ with average pay). Heatmaps of the joint inflation–wage distribution display smooth gradients without threshold effects or interaction patterns. The negative δ_2, the residual association between local inflation and Republican vote gains conditional on real wage growth, is therefore not an artifact of the linear functional form. The negative relationship is robust to flexible specifications, ruling out hidden reversals, threshold effects, or non-monotonic patterns that could reconcile the result with a standard incumbent-punishment prediction. The non-linearity tests also confirm that the null turnout–real-wage-growth association under average annual pay (Appendix H) is not a specification artifact.

Table H.1
Baseline Specification: Inflation, Wages, and Presidential Electoral Outcomes (Average Annual Pay)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Two-Party Republican Vote Share Change, 2020–2024								
Inflation	−0.0845	0.0137	−0.1542	−0.0234	−0.0573	−0.0603	−0.0085	−0.0526
	(0.0242)***	(0.0187)	(0.0535)***	(0.0164)	(0.0281)**	(0.0233)***	(0.0165)	(0.0258)**
	[0.0295]***	[0.0217]	[0.0527]***	[0.0178]	[0.0379]	[0.0293]**	[0.0232]	[0.0311]*
Nominal Wage Growth	−0.0255	−0.0261	−0.0190	−0.0257	−0.0259	−0.0280	−0.0258	−0.0260
	(0.0143)*	(0.0145)*	(0.0136)	(0.0145)*	(0.0144)*	(0.0146)*	(0.0145)*	(0.0144)*
	[0.0171]	[0.0169]	[0.0167]	[0.0169]	[0.0172]	[0.0167]*	[0.0170]	[0.0171]
Observations	3,102	3,102	3,102	3,102	3,102	3,102	3,102	3,102
R^2	0.412	0.409	0.420	0.409	0.410	0.410	0.409	0.410
Panel B: Republican Vote Margin Change, 2020–2024								
Inflation	−0.0771	0.0196	−0.1229	−0.0248	−0.0457	−0.0740	−0.0052	−0.0528
	(0.0226)***	(0.0180)	(0.0528)**	(0.0159)	(0.0267)*	(0.0248)***	(0.0155)	(0.0242)**
	[0.0279]***	[0.0224]	[0.0546]**	[0.0173]	[0.0339]	[0.0296]**	[0.0225]	[0.0296]*
Nominal Wage Growth	−0.0256	−0.0262	−0.0205	−0.0257	−0.0260	−0.0286	−0.0260	−0.0261
	(0.0146)*	(0.0147)*	(0.0147)	(0.0147)*	(0.0147)*	(0.0149)*	(0.0147)*	(0.0147)*
	[0.0132]*	[0.0130]**	[0.0130]	[0.0131]*	[0.0132]**	[0.0129]**	[0.0132]**	[0.0132]**
Observations	3,102	3,102	3,102	3,102	3,102	3,102	3,102	3,102
R^2	0.369	0.367	0.373	0.367	0.367	0.368	0.366	0.367
Panel C: Turnout Change, 2020–2024								
Inflation	0.0021	−0.0047	−0.0087	0.0775	−0.0494	0.0331	0.0538	−0.0272
	(0.0179)	(0.0155)	(0.0220)	(0.0178)***	(0.0201)**	(0.0211)	(0.0167)***	(0.0210)
	[0.0320]	[0.0175]	[0.0372]	[0.0212]***	[0.0268]*	[0.0345]	[0.0199]***	[0.0338]
Nominal Wage Growth	−0.0024	−0.0024	−0.0020	−0.0036	−0.0023	−0.0013	−0.0034	−0.0024
	(0.0193)	(0.0193)	(0.0195)	(0.0191)	(0.0192)	(0.0192)	(0.0192)	(0.0192)
	[0.0148]	[0.0148]	[0.0153]	[0.0149]	[0.0148]	[0.0147]	[0.0147]	[0.0148]
Observations	3,102	3,102	3,102	3,102	3,102	3,102	3,102	3,102
R^2	0.397	0.397	0.397	0.401	0.398	0.398	0.400	0.398
Price measure	Total	Food	Transport	Housing	Childcare	Healthcare	Other	Taxes
Wage measure	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Observations are at the county level. Each column reports estimates of specification (1) for a different price category k. Wage measure is average annual pay (QCEW) throughout. All variables standardized to zero mean and unit variance; coefficients are in standard-deviation units. Robust standard errors in parentheses; cluster-bootstrapped standard errors (100 replications, clustered at the state level) in brackets. All specifications include state fixed effects and controls for urbanization share, population density, median age, educational attainment (three shares), real GDP, and demographic composition. Significance stars reported next to each standard error type. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

H. Robustness: Average Annual Pay

Tables H.1–H.3 replicate the three main regression tables using average annual pay (QCEW) instead of total covered wages as the wage measure. Average annual pay isolates changes in per-worker compensation, whereas total covered wages also reflect changes in employment levels and industrial composition.

For partisan outcomes (vote share and vote margin), the results are qualitatively robust to this alternative wage measure. Inflation coefficients are nearly identical in magnitude and significance across the two measures; the wage–price gap remains negatively associated with Republican gains in the re-parameterized specification; and prices retain a strong negative residual association conditional on the gap ($\delta_2 \neq 0$). In the gap-only model, real wage growth alone continues to explain little of the variation in partisan vote swings.

For turnout, however, the results diverge. The positive association between the wage–price gap and turnout growth documented in the main tables (Tables 3–4) is absent when wages are measured by average annual pay: the gap coefficient is close to zero and statistically insignificant across most price categories. This suggests that the turnout–real-wage-growth channel operates primarily through aggregate labor market strength, reflecting both higher per-worker compensation and employment growth, rather than through per-worker pay gains alone. Appendix G.4 confirms that this null result is not an artifact of the linear functional form: binscatter plots, spline regressions, and heatmaps using average pay show no hidden non-linear turnout–gap relationship.

I. Congressional District Analysis

This appendix reports the full congressional district (CD) analysis summarized in Section 6.4. Tables I.2–I.4 mirror the county-level Tables 2–4: the baseline specification, the wage–price gap re-parameterization, and the restricted gap-only model, for four House electoral outcomes (change in two-party Republican vote share, change in the Republican vote margin, change in turnout, and change in challenger-party vote share, each between 2020 and 2024). The challenger

Table H.2
Wage–Price Gap Re-parameterization: Real Wage Growth, Prices, and Presidential Electoral Outcomes (Average Annual Pay)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Two-Party Republican Vote Share Change, 2020–2024								
Real Wage Growth	−0.0366	−0.0375	−0.0260	−0.0364	−0.0369	−0.0407	−0.0369	−0.0371
	(0.0205)*	(0.0208)*	(0.0187)	(0.0206)*	(0.0204)*	(0.0211)*	(0.0207)*	(0.0206)*
	[0.0245]	[0.0243]	[0.0229]	[0.0240]	[0.0244]	[0.0242]*	[0.0243]	[0.0244]
Inflation	−0.1100	−0.0124	−0.1732	−0.0491	−0.0832	−0.0883	−0.0344	−0.0786
	(0.0285)***	(0.0230)	(0.0544)***	(0.0215)**	(0.0314)***	(0.0283)***	(0.0214)	(0.0300)***
	[0.0352]***	[0.0259]	[0.0572]***	[0.0251]*	[0.0426]*	[0.0305]***	[0.0284]	[0.0341]**
Observations	3,102	3,102	3,102	3,102	3,102	3,102	3,102	3,102
R^2	0.412	0.409	0.420	0.409	0.410	0.410	0.409	0.410
Panel B: Republican Vote Margin Change, 2020–2024								
Real Wage Growth	−0.0368	−0.0377	−0.0281	−0.0365	−0.0370	−0.0415	−0.0372	−0.0373
	(0.0210)*	(0.0211)*	(0.0201)	(0.0209)*	(0.0209)*	(0.0216)*	(0.0211)*	(0.0210)*
	[0.0190]*	[0.0187]**	[0.0178]	[0.0187]*	[0.0189]**	[0.0187]**	[0.0188]**	[0.0189]**
Inflation	−0.1027	−0.0066	−0.1434	−0.0505	−0.0717	−0.1026	−0.0313	−0.0789
	(0.0266)***	(0.0215)	(0.0517)***	(0.0226)**	(0.0293)**	(0.0315)***	(0.0206)	(0.0288)***
	[0.0299]***	[0.0250]	[0.0562]**	[0.0223]**	[0.0355]**	[0.0302]***	[0.0256]	[0.0298]***
Observations	3,102	3,102	3,102	3,102	3,102	3,102	3,102	3,102
R^2	0.369	0.367	0.373	0.367	0.367	0.368	0.366	0.367
Panel C: Turnout Change, 2020–2024								
Real Wage Growth	−0.0035	−0.0034	−0.0027	−0.0050	−0.0033	−0.0019	−0.0049	−0.0034
	(0.0277)	(0.0277)	(0.0267)	(0.0271)	(0.0273)	(0.0278)	(0.0274)	(0.0274)
	[0.0213]	[0.0213]	[0.0209]	[0.0211]	[0.0211]	[0.0213]	[0.0210]	[0.0212]
Inflation	−0.0003	−0.0070	−0.0107	0.0739	−0.0517	0.0318	0.0504	−0.0296
	(0.0243)	(0.0233)	(0.0247)	(0.0273)***	(0.0274)*	(0.0278)	(0.0252)**	(0.0264)
	[0.0329]	[0.0238]	[0.0348]	[0.0249]***	[0.0283]*	[0.0382]	[0.0242]**	[0.0336]
Observations	3,102	3,102	3,102	3,102	3,102	3,102	3,102	3,102
R^2	0.397	0.397	0.397	0.401	0.398	0.398	0.400	0.398
Price measure	Total	Food	Transport	Housing	Childcare	Healthcare	Other	Taxes
Wage measure	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Observations are at the county level. Each column reports estimates of specification (3) for a different price category k. $G_i^{(k)} \equiv \Delta W_i - \Delta P_i^{(k)}$ is category-specific real wage growth. The gap and all other variables are standardized to zero mean and unit variance. This specification is algebraically equivalent to the baseline (Table H.1); R^2 values are identical by construction. Wage measure is average annual pay (QCEW) throughout. Robust standard errors in parentheses; cluster-bootstrapped standard errors (100 replications, clustered at the state level) in brackets. All specifications include state fixed effects and the full control set. Significance stars reported next to each standard error type. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

is defined as the party that did not win the district in the prior House cycle (2018 for 2020, 2022 for 2024). Table I.1 reports summary statistics for the CD estimation sample.

Baseline specification

The CD-level estimates for inflation and wages in the baseline specification are substantially noisier than the county results. Point estimates on inflation are generally negative, consistent with the presidential-level results, but magnitudes are smaller and precision is markedly lower, particularly under state-clustered inference. Wage coefficients are also close to zero and typically statistically insignificant. This attenuation reflects the smaller effective sample ($N \approx 1{,}270$ district-congress observations versus 3,102 counties) and the aggregation of county-level covariates to district boundaries, which introduces additional measurement error.

For turnout, the CD baseline results diverge sharply from the county-level pattern. Nominal wage growth enters with a *negative* sign and is precisely estimated: the coefficient ranges from −0.222 to −0.266 across price categories, significant at the five percent level under state-clustered inference in every column. This contrasts with the county-level presidential analysis, where nominal wage growth is positively associated with turnout.

For challenger-party vote share the estimates reveal a mixed picture: while some inflation categories are associated with modestly higher challenger vote shares, the estimates are imprecise and do not consistently reach statistical significance. Local inflation shocks, as measured here, do not reliably translate into anti-incumbent vote switching in House races.

Wage–price gap re-parameterization

Table I.3 reports the re-parameterized specification that decomposes the economic environment into the wage–price gap $G_i^{(k)}$ and nominal price growth $\Delta P_i^{(k)}$, mirroring the county-level Table 3.

Table H.3
Restricted Model: Real Wage Growth Only, Prices Omitted (Average Annual Pay)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Two-Party Republican Vote Share Change, 2020–2024								
Real Wage Growth	0.0174	−0.0291	0.0532	−0.0073	−0.0011	−0.0044	−0.0145	−0.0026
	(0.0179)	(0.0172)*	(0.0333)	(0.0157)	(0.0189)	(0.0173)	(0.0161)	(0.0179)
	[0.0236]	[0.0213]	[0.0356]	[0.0182]	[0.0232]	[0.0253]	[0.0219]	[0.0248]
Observations	3,102	3,102	3,102	3,102	3,102	3,102	3,102	3,102
R^2	0.408	0.409	0.410	0.408	0.408	0.408	0.408	0.408
Panel B: Republican Vote Margin Change, 2020–2024								
Real Wage Growth	0.0136	−0.0332	0.0375	−0.0065	−0.0062	0.0006	−0.0168	−0.0027
	(0.0180)	(0.0177)*	(0.0338)	(0.0147)	(0.0192)	(0.0165)	(0.0158)	(0.0176)
	[0.0211]	[0.0194]*	[0.0339]	[0.0152]	[0.0202]	[0.0218]	[0.0195]	[0.0221]
Observations	3,102	3,102	3,102	3,102	3,102	3,102	3,102	3,102
R^2	0.366	0.367	0.367	0.366	0.366	0.366	0.366	0.366
Panel C: Turnout Change, 2020–2024								
Real Wage Growth	−0.0033	0.0014	0.0021	−0.0489	0.0190	−0.0149	−0.0377	0.0096
	(0.0214)	(0.0189)	(0.0228)	(0.0186)***	(0.0209)	(0.0223)	(0.0188)**	(0.0222)
	[0.0229]	[0.0154]	[0.0251]	[0.0170]***	[0.0199]	[0.0214]	[0.0164]**	[0.0234]
Observations	3,102	3,102	3,102	3,102	3,102	3,102	3,102	3,102
R^2	0.397	0.397	0.397	0.399	0.398	0.398	0.399	0.397
Price measure	Total	Food	Transport	Housing	Childcare	Healthcare	Other	Taxes
Wage measure	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay	Average Pay
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Observations are at the county level. Each column reports estimates of the restricted specification (4), which includes the wage–price gap $G_i^{(k)} \equiv \Delta W_i - \Delta P_i^{(k)}$ but omits $\Delta P_i^{(k)}$ as a separate regressor. The gap is standardized to zero mean and unit variance. Wage measure is average annual pay (QCEW) throughout. Robust standard errors in parentheses; cluster-bootstrapped standard errors (100 replications, clustered at the state level) in brackets. All specifications include state fixed effects and the full control set. Significance stars reported next to each standard error type. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

For Republican vote share and margin, the gap coefficient δ_1 is close to zero and statistically insignificant across all price categories. For total inflation, the estimates are −0.014 for vote-share change and −0.012 for margin change, an order of magnitude smaller than the corresponding county-level estimates (−0.052 and −0.048). The residual price coefficient δ_2 is negative for several categories—notably healthcare (−0.174) and taxes (−0.151)—but these estimates reach significance only under robust standard errors and lose precision under state clustering. The direction is consistent with the county-level pattern, but the evidence is insufficient to reject $H_0^{(2)}$ at the district level for partisan outcomes.

For turnout, the results reveal a striking departure from the county-level pattern. The gap coefficient is *negative* and precisely estimated across all eight price categories: for total inflation, $\delta_1 = -0.272$ ($p < 0.05$ under state-clustered inference), and estimates range from −0.333 (childcare) to −0.244 (housing), all significant under both robust and clustered standard errors. This sign reversal—positive at the county level for presidential turnout (Table 3, Panel C), negative at the district level for House turnout—is the most notable difference between the presidential and congressional analyses.

Several factors may account for the divergence. First, the two analyses measure different electoral margins: county regressions capture presidential turnout, whereas the CD regressions capture House turnout. Although both are measured in presidential election years, House-specific turnout dynamics—driven by district competitiveness, incumbency advantage, and candidate salience—may respond differently to local economic conditions than presidential turnout. Second, the direction of the reversal is consistent with the paper's broader finding that economically dynamic areas have trended Democratic: districts with stronger real wage growth tend to be safe Democratic seats where House races are less competitive, dampening district-level turnout mobilization for the congressional contest even as presidential turnout in these same areas may grow. Third, the aggregation from counties to districts mechanically smooths within-state variation and may introduce composition effects that shift the estimated relationship. The negative coefficient is therefore best interpreted as reflecting the distinct electoral context of House races rather than contradicting the county-level mobilization channel for presidential elections.

The residual price coefficient δ_2 is also negative and statistically significant for turnout in most columns—for instance, −0.258 for total inflation, −0.346 for childcare, and −0.295 for healthcare, all significant under state clustering. Combined with the negative δ_1, the CD turnout results indicate that both dimensions of local economic dynamism—real wage growth and residual nominal activity—are associated with *lower* House turnout growth, the opposite of the presidential-turnout pattern. This divergence underscores that the mechanisms linking local economic conditions to participation may operate differently across levels of government.

Table I.1
Summary Statistics — Congressional District Level

Variable	Mean	SD	Min	P25	P50	P75	Max	N
Panel A: Electoral Outcomes — Changes								
Change in Republican vote share pct (alt) 2024 vs 2020 (pp)	2.7725	14.5111	-54.4782	-2.0228	1.6257	6.9168	76.7762	1270
Change in Democratic vote share pct (alt) 2024 vs 2020 (pp)	-2.7725	14.5111	-76.7762	-6.9168	-1.6257	2.0228	54.4782	1270
Change in turnout, 2020–2024 (pp)	-1.1979	8.9272	-58.6538	-4.6976	-0.7063	2.8496	100.0004	1257
Change in Rep. vote margin, 2020–2024 (pp)	5.2059	26.2494	-108.9554	-4.0025	3.6392	13.4145	153.5524	1270
Change in challenger two-party vote share, 2020–2024 (pp)	-0.3500	13.6021	-46.5021	-5.0370	-0.1295	4.3564	45.5218	1270
Panel B: Inflation Changes 2021–2024								
Inflation food 2021–2024 (%) (CD agg)	28.3555	5.6174	4.7114	24.5436	28.0402	31.5192	63.0183	1270
Inflation transportation 2021–2024 (%) (CD agg)	27.5435	9.1898	12.9496	23.1298	27.4612	29.8338	139.7267	1270
Inflation housing 2021–2024 (%) (CD agg)	29.1304	9.4612	-6.0848	22.9856	28.7582	34.7958	59.2740	1270
Inflation childcare 2021–2024 (%) (CD agg)	27.2970	32.0684	-28.6451	5.3153	22.4954	45.1684	179.0404	1270
Inflation healthcare 2021–2024 (%) (CD agg)	17.3690	10.5466	-12.1624	11.7117	17.5795	22.6908	89.8601	1270
Inflation othernecessities 2021–2024 (%) (CD agg)	19.6187	5.9073	-6.8493	16.0424	19.6606	23.3562	37.6495	1270
Inflation taxes 2021–2024 (%) (CD agg)	52.1004	19.4494	10.9503	35.8228	52.4396	66.3782	115.9399	1270
Inflation total 2021–2024 (%) (CD agg)	27.0154	7.6776	9.4235	21.0616	27.0834	32.0112	54.8455	1270
Panel C: Nominal Wage Growth								
Change in total wages, 2021–2024 (%)	19.6231	6.0939	-99.9966	15.7724	19.5875	22.9625	39.6619	1270
Change in average pay, 2021–2024 (%)	12.0942	3.0315	-5.0221	10.0136	12.5690	14.5290	19.8954	1270
Panel D: Real Wage Growth (Total Wages – Inflation)								
Real wage growth: total (Total Wages)	-7.3924	9.6474	-123.6485	-13.0591	-6.5442	-1.1732	17.0532	1270
Real wage growth: food (Total Wages)	-8.7325	8.4608	-130.7533	-13.4156	-7.7972	-3.6932	13.0914	1270
Real wage growth: transportation (Total Wages)	-7.9204	11.4246	-128.2221	-11.7332	-7.1896	-2.6737	11.5791	1270
Real wage growth: housing (Total Wages)	-9.5073	9.1512	-122.5997	-14.5879	-8.6946	-4.3744	20.4806	1270
Real wage growth: childcare (Total Wages)	-7.6739	33.4083	-168.0768	-22.5977	-1.1342	15.0679	49.7485	1270
Real wage growth: healthcare (Total Wages)	2.2541	11.8431	-112.2134	-4.7921	2.2039	9.0528	37.7086	1270
Real wage growth: othernecessities (Total Wages)	0.0044	6.9654	-117.1529	-3.5652	0.1121	3.6635	25.2430	1270
Real wage growth: taxes (Total Wages)	-32.4773	19.7202	-142.2719	-46.2731	-31.9774	-17.0580	8.9568	1270
Panel E: Controls								
Percent of 2020 land classified as urban (CD agg)	0.3360	0.2844	0.0043	0.1070	0.2364	0.5157	1.0000	1270
2020 population density of the county (per square mile) (CD agg)	2378.3489	7143.2386	12.7143	287.0554	704.5940	1787.8507	7.48e+04	1270
Median age for total population (CD agg)	38.7762	3.0334	28.5059	37.0853	38.4216	40.3455	55.9452	1270
Percent of adults who are high school graduates (or equivalent), 2019-23 (CD agg	26.1448	5.8529	9.9035	22.1227	26.1720	29.8108	43.6617	1270
Percent of adults completing some college or associate degree, 2019-23 (CD agg)	28.1863	4.2884	13.3611	25.5262	28.5886	31.2454	38.3545	1270
Percent of adults with a bachelor's degree or higher, 2019-23 (CD agg)	34.9019	9.1020	15.6038	28.1380	33.9776	41.0101	69.2567	1270
Real GDP (thousands of chained 2017 dollars) (2020) (CD agg)	4.65e+07	3.39e+07	968.3064	3.08e+07	3.90e+07	5.19e+07	3.38e+08	1270
Percent: HISPANIC OR LATINO: Total population: Hispanic or Latino (of any race)	18.8180	15.9063	1.2685	6.6860	12.9521	26.6364	89.0231	1270
Percent: RACE: Total population: One Race: White (%) (DP1 2020) (CD agg)	61.3989	17.4949	14.1000	47.0998	61.8209	76.0666	94.3410	1270

Notes: Observations are congressional district by congressional cycle in the CD estimation sample ($N \approx 1{,}270$ district-cycle observations across the 116th, 117th, and 118th Congresses). Panel A reports 2020–2024 changes in district-level U.S. House electoral outcomes computed from MIT Election Lab House returns. Panel B reports 2021–2024 percent changes in the Economic Policy Institute Family Budget Calculator costs aggregated from county to district by population-weighted average. Panel C reports 2021–2024 percent changes in total covered wages and average annual pay (QCEW), aggregated to the district. Panel D reports the category-specific wage–price gap.

For challenger-party vote share, the gap coefficient is small and positive (ranging from 0.013 to 0.030) and statistically insignificant across all columns. The residual price coefficient is likewise close to zero, with the exception of transportation (0.150, significant under robust but not clustered inference).

Gap-only specification

Table I.4 estimates the restricted model that includes only the wage–price gap, omitting prices as a separate regressor. For partisan outcomes (Panels A and B), the gap coefficient is small and generally insignificant, consistent with the county-level finding that real wage growth alone explains little variation in partisan swings. The same omitted-variable logic identified at the county level applies here: the negative residual price coefficients in Table I.3 combined with the negative correlation between the gap and inflation generate a positive bias that pushes the already-small unrestricted estimates toward zero or slightly positive values (e.g., 0.030 for total inflation and 0.107 for healthcare in Panel A). For turnout (Panel C), the gap coefficient remains negative—ranging from −0.041 to −0.166—and is statistically distinguishable from zero under state clustering for food (−0.166), housing (−0.077), and other necessities (−0.110), confirming that the negative House-turnout association documented in Table I.3 is not an artifact of including prices in the specification. The magnitudes are smaller than in the unrestricted model (roughly −0.04 to −0.17 versus −0.24 to −0.33); by the same omitted-variable reasoning, dropping the negatively signed δ_2 biases the gap coefficient toward zero, attenuating the negative association. For the challenger-party outcome (Panel D), the gap coefficient is close to zero throughout, reinforcing the absence of an incumbent-punishment channel in House elections.

J. A comment on Flavin (2026) results

Flavin (2026) examines the relationship between county-level inflation and electoral outcomes in the 2024 U.S. presidential election. Using EPI Family Budget Calculator data, he measures inflation over 2023–2024 across eight expenditure categories (total, housing, food, transportation, healthcare, other necessities, childcare, and taxes) and

Table I.2
Congressional District: Inflation, Wages, and House Electoral Outcomes

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Two-Party Republican Vote Share Change, 2020–2024								
Inflation	−0.0697	−0.0068	0.0142	0.0127	0.0290	−0.1699	0.0136	−0.1383
	(0.0611)	(0.0398)	(0.0566)	(0.0578)	(0.0639)	(0.0669)**	(0.0542)	(0.0705)*
	[0.0923]	[0.0670]	[0.1592]	[0.0607]	[0.1245]	[0.1472]	[0.0512]	[0.0977]
Nominal Wage Growth	−0.0117	−0.0062	−0.0053	−0.0094	0.0012	−0.0039	−0.0093	−0.0123
	(0.0420)	(0.0425)	(0.0425)	(0.0434)	(0.0427)	(0.0425)	(0.0427)	(0.0417)
	[0.0754]	[0.0778]	[0.0797]	[0.0800]	[0.0749]	[0.0775]	[0.0794]	[0.0741]
Observations	1,270	1,270	1,270	1,270	1,270	1,270	1,270	1,270
R^2	0.108	0.108	0.108	0.108	0.108	0.115	0.108	0.111
Panel B: Republican Vote Margin Change, 2020–2024								
Inflation	−0.0643	−0.0073	0.0338	0.0174	0.0278	−0.1790	0.0200	−0.1401
	(0.0621)	(0.0404)	(0.0546)	(0.0607)	(0.0675)	(0.0708)**	(0.0572)	(0.0723)*
	[0.0811]	[0.0749]	[0.1477]	[0.0639]	[0.1121]	[0.1407]	[0.0563]	[0.0946]
Nominal Wage Growth	−0.0103	−0.0053	−0.0035	−0.0096	0.0018	−0.0025	−0.0098	−0.0114
	(0.0417)	(0.0426)	(0.0424)	(0.0437)	(0.0419)	(0.0423)	(0.0428)	(0.0416)
	[0.0768]	[0.0811]	[0.0796]	[0.0838]	[0.0757]	[0.0752]	[0.0840]	[0.0732]
Observations	1,294	1,294	1,294	1,294	1,294	1,294	1,294	1,294
R^2	0.072	0.071	0.071	0.071	0.071	0.079	0.071	0.075
Panel C: Turnout Change, 2020–2024								
Inflation	−0.0287	0.0956	0.0411	0.0448	−0.0847	−0.0656	0.0797	−0.0013
	(0.0523)	(0.0326)***	(0.0424)	(0.0391)	(0.0596)	(0.0373)*	(0.0371)**	(0.0567)
	[0.0600]	[0.0605]	[0.1288]	[0.0415]	[0.0694]	[0.0519]	[0.0357]**	[0.0659]
Nominal Wage Growth	−0.2291	−0.2238	−0.2226	−0.2491	−0.2609	−0.2293	−0.2655	−0.2263
	(0.0800)***	(0.0796)***	(0.0804)***	(0.0840)***	(0.0863)***	(0.0796)***	(0.0828)***	(0.0798)***
	[0.1006]**	[0.1044]**	[0.1035]**	[0.1116]**	[0.1032]**	[0.1027]**	[0.1115]**	[0.1011]**
Observations	1,257	1,257	1,257	1,257	1,257	1,257	1,257	1,257
R^2	0.236	0.239	0.236	0.237	0.237	0.237	0.238	0.236
Panel D: Challenger Two-Party Vote Share Change, 2020–2024								
Inflation	−0.0147	0.0226	0.1322	−0.0687	0.0148	0.0116	−0.0538	0.0256
	(0.0596)	(0.0410)	(0.0563)**	(0.0433)	(0.0617)	(0.0488)	(0.0417)	(0.0649)
	[0.1026]	[0.0854]	[0.1711]	[0.0685]	[0.1467]	[0.0919]	[0.0587]	[0.1176]
Nominal Wage Growth	0.0112	0.0135	0.0177	0.0310	0.0160	0.0123	0.0258	0.0136
	(0.0401)	(0.0396)	(0.0404)	(0.0411)	(0.0405)	(0.0396)	(0.0404)	(0.0398)
	[0.1151]	[0.1135]	[0.1121]	[0.1204]	[0.1225]	[0.1133]	[0.1171]	[0.1141]
Observations	1,270	1,270	1,270	1,270	1,270	1,270	1,270	1,270
R^2	0.148	0.148	0.151	0.149	0.148	0.148	0.149	0.148
Price measure	Total	Food	Transport	Housing	Childcare	Healthcare	Other	Taxes
Wage measure	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Observations are at the congressional district level. Each column reports estimates of specification (1) for a different price category k, using U.S. House election outcomes. Wage measure is total covered wages (QCEW) throughout. All variables standardized to zero mean and unit variance. Robust standard errors in parentheses; cluster-bootstrapped standard errors (100 replications, clustered at the state level) in brackets. All specifications include state fixed effects and the full control set. The challenger outcome is defined as the two-party vote share of the party that did not win the prior House cycle (2018 for 2020, 2022 for 2024). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

regresses the change in Republican two-party vote share on each standardized inflation category in turn, controlling for the county unemployment rate in 2023, bachelor's degree share, median age, evangelical congregation rate, percent white, and log population. He reports a positive and statistically significant association for several categories, notably total, childcare, and tax inflation, and concludes that rising prices contributed to Republican vote gains.

This appendix reassesses the robustness of that finding. Several aspects of Flavin's empirical specification differ from ours and may account for the divergence in results. First, his specification omits wage data and therefore cannot distinguish between the real wage growth channel and nominal price salience, the central decomposition underlying our wage–price gap reparameterization. Second, his baseline model does not include state fixed effects, leaving the inflation coefficient identified in part from cross-state variation that is correlated with long-standing partisan alignments. Third, and most importantly, as we demonstrate below, the reported statistical inference is not robust: once heteroskedasticity-consistent standard errors or clustered error structures are employed, the estimated effects are no longer statistically significant.

We organize the analysis around three main findings. First, we replicate Flavin's specification exactly and show that the reported significance does not survive standard inference corrections. Second, we demonstrate that the positive sign of the inflation coefficient is an artifact of between-state confounding, which reverses once state fixed effects are included. Third, we show that omitting wage controls prevents distinguishing nominal price salience from real income

Table I.3
Congressional District: Wage–Price Gap Re-parameterization (Most Demanding)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Two-Party Republican Vote Share Change, 2020–2024								
Real Wage Growth	−0.0140	−0.0077	−0.0067	−0.0092	0.0015	−0.0046	−0.0094	−0.0141
	(0.0500)	(0.0524)	(0.0532)	(0.0425)	(0.0544)	(0.0499)	(0.0432)	(0.0477)
	[0.0897]	[0.0959]	[0.0998]	[0.0785]	[0.0956]	[0.0909]	[0.0804]	[0.0847]
Inflation	−0.0815	−0.0130	0.0089	0.0034	0.0302	−0.1738	0.0043	−0.1506
	(0.0759)	(0.0627)	(0.0751)	(0.0633)	(0.0840)	(0.0764)**	(0.0628)	(0.0824)*
	[0.1139]	[0.1153]	[0.1767]	[0.1036]	[0.1486]	[0.1574]	[0.1010]	[0.1172]
Observations	1,270	1,270	1,270	1,270	1,270	1,270	1,270	1,270
R^2	0.108	0.108	0.108	0.108	0.108	0.115	0.108	0.111
Panel B: Republican Vote Margin Change, 2020–2024								
Real Wage Growth	−0.0122	−0.0065	−0.0044	−0.0094	0.0023	−0.0029	−0.0100	−0.0131
	(0.0496)	(0.0526)	(0.0531)	(0.0429)	(0.0535)	(0.0496)	(0.0433)	(0.0476)
	[0.0913]	[0.1000]	[0.0997]	[0.0821]	[0.0967]	[0.0882]	[0.0851]	[0.0837]
Inflation	−0.0745	−0.0125	0.0303	0.0078	0.0297	−0.1815	0.0102	−0.1516
	(0.0742)	(0.0641)	(0.0713)	(0.0646)	(0.0852)	(0.0795)**	(0.0645)	(0.0820)*
	[0.0929]	[0.1288]	[0.1628]	[0.1039]	[0.1248]	[0.1504]	[0.1085]	[0.1105]
Observations	1,294	1,294	1,294	1,294	1,294	1,294	1,294	1,294
R^2	0.072	0.071	0.071	0.071	0.071	0.079	0.071	0.075
Panel C: Turnout Change, 2020–2024								
Real Wage Growth	−0.2724	−0.2760	−0.2789	−0.2442	−0.3330	−0.2690	−0.2687	−0.2587
	(0.0951)***	(0.0982)***	(0.1008)***	(0.0823)***	(0.1102)***	(0.0934)***	(0.0838)***	(0.0912)***
	[0.1196]**	[0.1288]**	[0.1297]**	[0.1094]**	[0.1317]**	[0.1205]**	[0.1129]**	[0.1156]**
Inflation	−0.2578	−0.1282	−0.1815	−0.2042	−0.3457	−0.2949	−0.1857	−0.2276
	(0.0983)***	(0.0907)	(0.0984)*	(0.0801)**	(0.1234)***	(0.0836)***	(0.0804)**	(0.0979)**
	[0.1058]**	[0.1367]	[0.1364]	[0.1010]**	[0.1254]***	[0.1252]**	[0.1086]*	[0.1043]**
Observations	1,257	1,257	1,257	1,257	1,257	1,257	1,257	1,257
R^2	0.236	0.239	0.236	0.237	0.237	0.237	0.238	0.236
Panel D: Challenger Two-Party Vote Share Change, 2020–2024								
Real Wage Growth	0.0133	0.0166	0.0221	0.0304	0.0204	0.0144	0.0261	0.0155
	(0.0477)	(0.0488)	(0.0507)	(0.0403)	(0.0517)	(0.0464)	(0.0409)	(0.0455)
	[0.1368]	[0.1400]	[0.1405]	[0.1181]	[0.1564]	[0.1329]	[0.1185]	[0.1305]
Inflation	−0.0035	0.0361	0.1499	−0.0377	0.0308	0.0239	−0.0280	0.0391
	(0.0772)	(0.0585)	(0.0744)**	(0.0576)	(0.0814)	(0.0626)	(0.0568)	(0.0805)
	[0.1747]	[0.1548]	[0.2047]	[0.1261]	[0.2202]	[0.1400]	[0.1283]	[0.1762]
Observations	1,270	1,270	1,270	1,270	1,270	1,270	1,270	1,270
R^2	0.148	0.148	0.151	0.149	0.148	0.148	0.149	0.148
Price measure	Total	Food	Transport	Housing	Childcare	Healthcare	Other	Taxes
Wage measure	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Observations are at the congressional district level. Each column reports estimates of specification (3) for a different price category k, using U.S. House election outcomes. $G_i^{(k)} \equiv \Delta W_i - \Delta P_i^{(k)}$ is category-specific real wage growth. Wage measure is total covered wages (QCEW) throughout. All variables standardized. Robust standard errors in parentheses; cluster-bootstrapped standard errors (100 replications, clustered at the state level) in brackets. All specifications include state fixed effects and the full control set. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

effects, and that the inflation coefficient is negative regardless of which demographic control set is employed once wages are included. We conclude with a discussion of additional specification concerns.

Replication and statistical inference

We load Flavin's replication data and code and reproduce his main results (corresponding to Figure 1 in his paper). His specification regresses the change in Republican two-party vote share on each of eight standardized EPI Family Budget Calculator inflation categories (one at a time), controlling for the county unemployment rate in 2023, bachelor's degree share, median age, evangelical congregation rate, percent white, and log population. Table J.1, Panel A confirms his reported coefficients and significance levels under conventional OLS standard errors: total inflation (0.068, $p < 0.01$), childcare (0.148, $p < 0.01$), and taxes (0.093, $p < 0.01$) are positively associated with gains in Trump's vote share, while healthcare inflation (−0.106, $p < 0.01$) carries a negative sign. The remaining four categories (housing, food, transportation, and other necessities) are not statistically significant even under OLS standard errors. We replicate these results exactly.

The key issue is inference. Panel A also reports heteroskedasticity-robust and state-clustered standard errors for the same point estimates, that is, no changes to the model, only to the standard errors. EPI inflation draws on underlying data sources (HUD Fair Market Rents, USDA food plans, ACA premiums) that share common metro-area or regional components, inducing spatial correlation among counties within the same state. Clustering at the state level is therefore the appropriate error structure; failing to do so understates sampling variability.

Table I.4
Congressional District: Restricted Model, Real Wage Growth Only

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Two-Party Republican Vote Share Change, 2020–2024								
Real Wage Growth	0.0298	0.0012	−0.0114	−0.0113	−0.0171	0.1000	−0.0123	0.0595
	(0.0420)	(0.0329)	(0.0402)	(0.0429)	(0.0435)	(0.0525)*	(0.0413)	(0.0459)
	[0.0640]	[0.0502]	[0.0872]	[0.0412]	[0.0687]	[0.0987]	[0.0364]	[0.0657]
Observations	1,270	1,270	1,270	1,270	1,270	1,270	1,270	1,270
R^2	0.108	0.108	0.108	0.108	0.108	0.111	0.108	0.109
Panel B: Republican Vote Margin Change, 2020–2024								
Real Wage Growth	0.0278	0.0021	−0.0206	−0.0143	−0.0159	0.1070	−0.0167	0.0611
	(0.0439)	(0.0326)	(0.0410)	(0.0451)	(0.0460)	(0.0553)*	(0.0435)	(0.0482)
	[0.0690]	[0.0488]	[0.0812]	[0.0447]	[0.0766]	[0.1000]	[0.0367]	[0.0670]
Observations	1,294	1,294	1,294	1,294	1,294	1,294	1,294	1,294
R^2	0.071	0.071	0.071	0.071	0.071	0.075	0.071	0.072
Panel C: Turnout Change, 2020–2024								
Real Wage Growth	−0.0899	−0.1656	−0.1472	−0.0765	−0.0670	−0.0407	−0.1104	−0.1080
	(0.0518)*	(0.0376)***	(0.0467)***	(0.0387)**	(0.0490)	(0.0455)	(0.0378)***	(0.0539)**
	[0.0698]	[0.0506]***	[0.0896]	[0.0442]*	[0.0882]	[0.0442]	[0.0405]***	[0.0740]
Observations	1,257	1,257	1,257	1,257	1,257	1,257	1,257	1,257
R^2	0.232	0.237	0.234	0.232	0.231	0.230	0.234	0.232
Panel D: Challenger Two-Party Vote Share Change, 2020–2024								
Real Wage Growth	0.0152	−0.0078	−0.0579	0.0543	0.0015	0.0000	0.0447	−0.0036
	(0.0377)	(0.0343)	(0.0393)	(0.0333)	(0.0408)	(0.0370)	(0.0325)	(0.0380)
	[0.0689]	[0.0667]	[0.0886]	[0.0567]	[0.0881]	[0.0815]	[0.0492]	[0.0758]
Observations	1,270	1,270	1,270	1,270	1,270	1,270	1,270	1,270
R^2	0.148	0.148	0.149	0.149	0.148	0.148	0.149	0.148
Price measure	Total	Food	Transport	Housing	Childcare	Healthcare	Other	Taxes
Wage measure	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages	Total Wages
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Observations are at the congressional district level. Each column reports estimates of specification (4), which includes the wage–price gap but omits prices as a separate regressor. U.S. House election outcomes. Wage measure is total covered wages (QCEW) throughout. All variables standardized. Robust standard errors in parentheses; cluster-bootstrapped standard errors (100 replications, clustered at the state level) in brackets. All specifications include state fixed effects and the full control set. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

With heteroskedasticity-robust standard errors, the changes are modest: standard errors increase by 5–15% and significance levels are largely preserved. With state-clustered standard errors, however, the picture changes dramatically. Standard errors increase by a factor of two to four; for example, the total inflation SE rises from 0.025 (OLS) to 0.086 (state-clustered), a 3.4-fold increase. Under state-clustered inference, only the childcare category retains statistical significance at the 5% level; the remaining seven categories, including total inflation, become statistically insignificant.

The role of state fixed effects

Flavin's baseline specification does not include state fixed effects. Without them, the inflation coefficient is identified from both within-state and between-state variation in prices. Between-state variation is problematic: cross-state differences in inflation regimes (persistently high-cost coastal states versus lower-cost interior states) are correlated with long-standing political alignments and secular partisan realignment trends. If higher-inflation states also experienced larger Republican gains for reasons unrelated to prices (differential mobilization, candidate-specific effects, compositional demographic change), then the between-state component of the inflation coefficient is confounded. State fixed effects restrict identification to within-state variation: whether, among counties in the *same state*, those facing larger local price increases shifted differentially toward one party.

Panel B of Table J.1 reports results from Flavin's specification augmented only with state fixed effects. The consequences are stark. Coefficients reverse sign for most categories: total inflation becomes −0.055 ($p < 0.10$), housing −0.057 ($p < 0.05$), and other necessities −0.071 ($p < 0.05$), all negative, directly contradicting the positive association reported in the original. Under cluster-bootstrapped standard errors (100 replications at the state level), none of the eight categories is statistically significant at conventional levels.

The sign reversal is itself diagnostic. It indicates that the positive OLS estimate was driven primarily by between-state covariation, that is, counties in high-inflation states that also happened to swing Republican, rather than by within-state differences in local price growth. Once this cross-state confound is absorbed, the within-state association between inflation and Republican gains turns negative, consistent with the results we report in the main text.

It is worth stating precisely what the two specifications identify and what they do not. Flavin's specification without state fixed effects identifies an inflation-vote association that pools within-state and between-state variation. This object includes, by construction, covariation between cross-state differences in cost-of-living and cross-state differences in long-run partisan trajectories. Our specification absorbs state means via fixed effects and therefore identifies a different object: the cross-county association between inflation and vote-share changes *among counties in the same state*, conditional on the wage–price gap and the baseline covariate set. This is the object most informative for the paper's causal question—whether differential local price experience within a state translates into differential local political behavior—but by design it is silent on whether the aggregate national or between-state pattern runs in the same direction. The divergence in sign between the two specifications is therefore not evidence that one is wrong. It indicates that the between-state and within-state components of the inflation-vote covariance move in opposite directions, a fact that is itself substantively informative about the geography of U.S. partisan realignment, and one that a specification pooling both sources of variation will obscure.

Omitted wage controls

Flavin's specification does not include any measure of nominal wage growth. This omission is consequential because local inflation and wage growth are positively correlated: counties with larger price increases tend also to experience larger nominal wage gains, partly through labor market adjustment and partly because both respond to common demand shocks. Without conditioning on wages, the inflation coefficient confounds two conceptually distinct channels: (i) a *nominal price salience* channel, whereby voters punish incumbents for rising prices regardless of purchasing power, and (ii) a *real income* channel, whereby voters respond to changes in living standards. Our baseline specification (1) separates these channels by including wage growth alongside inflation, and the wage–price gap reparameterization (3) provides a direct test of whether prices retain explanatory power *beyond* their effect on real wages.

To demonstrate that it is the inclusion of wages and state fixed effects, rather than differences in the demographic control set, that accounts for the sign divergence, we conduct a cross-specification exercise. We re-estimate our baseline on our data using Flavin's six control variables (unemployment rate, bachelor's degree share, median age, evangelical congregation rate, percent white, and log population) in place of our own controls, while retaining wages and state fixed effects.

Table J.2 reports the results. The inflation coefficient remains negative across all eight price categories: for total inflation the estimate is -0.066 ($p < 0.01$ with robust SEs); for other necessities -0.047 ($p < 0.05$); and for childcare -0.032 ($p < 0.05$). These estimates closely mirror our main results reported in Tables 2 and H.1. The choice of demographic controls does not drive the sign reversal; wages and state fixed effects do.

Additional specification concerns

Two of Flavin's controls raise endogeneity concerns. The county unemployment rate in 2023 is a "bad control" in the sense of Angrist and Pischke (2009): if inflation affects local labor market conditions (through real wage adjustments, changes in labor force participation, or firm entry and exit) then conditioning on unemployment absorbs part of the causal pathway from prices to electoral behavior, attenuating the estimated total effect. The evangelical congregation rate is a strong predictor of Republican vote-share *levels* but is nearly time-invariant; in a regression on vote-share *changes*, it functions more like a control for baseline partisanship than a demographic confounder, risking over-control that absorbs outcome variation attributable to economic conditions.[2]

Beyond the control set, two measurement choices are worth noting. Flavin measures inflation over 2023–2024 only, a single year of price changes, whereas we use cumulative 2021–2024 inflation spanning the full presidential term. A one-year window captures transitory price shocks rather than the persistent cost-of-living trajectory that voters plausibly attribute to incumbent economic management (Fiorina, 1981). Additionally, Flavin's election returns are sourced from a FOX News GitHub repository rather than from official state-certified counts; county-level figures from election-night tallies can diverge from certified results due to provisional ballots and late tabulation corrections, introducing measurement error in the dependent variable.

In sum, the positive association between county-level inflation and Republican vote gains reported in Flavin (2026) does not survive standard inference corrections or the inclusion of state fixed effects and wage controls. The reported

[2] In a panel first-differencing framework, a time-invariant regressor would be eliminated entirely. Its appearance in the cross-sectional regression is an artifact of the single-cross-section design.

Table J.1
Replication of Flavin (2026): Inference Sensitivity

	(1) Total	(2) Housing	(3) Food	(4) Transport.	(5) Healthcare	(6) Other Nec.	(7) Childcare	(8) Taxes
Panel A: Flavin (2026) specification (no state fixed effects)								
Inflation (std)	0.0682***	−0.0196	0.0338	0.0462*	−0.1060***	−0.0348	0.1482***	0.0925***
OLS	(0.0254)	(0.0267)	(0.0253)	(0.0279)	(0.0255)	(0.0273)	(0.0255)	(0.0254)
Robust	(0.0266)	(0.0271)	(0.0240)	(0.0263)	(0.0287)	(0.0286)	(0.0269)	(0.0276)
State-clustered	[0.0860]	[0.0573]	[0.0250]	[0.1041]	[0.0967]	[0.0612]	[0.0719]	[0.0873]
Observations	3,103	3,103	3,103	3,103	3,103	3,103	3,103	3,103
R^2	0.292	0.290	0.290	0.291	0.294	0.290	0.298	0.293
Panel B: Adding state fixed effects								
Inflation (std)	−0.0549*	−0.0574**	0.0394*	1.6133*	−0.0473	−0.0705**	0.0119	−0.0543*
Robust	(0.0301)	(0.0282)	(0.0214)	(0.8299)	(0.0298)	(0.0292)	(0.0343)	(0.0324)
Cluster-bootstrap	[0.0375]	[0.0261]	[0.0224]	[0.8549]	[0.0470]	[0.0326]	[0.0380]	[0.0393]
Observations	3,103	3,103	3,103	3,103	3,103	3,103	3,103	3,103
R^2	0.477	0.477	0.477	0.477	0.476	0.477	0.476	0.477
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Price measure	Total	Housing	Food	Transport.	Healthcare	Other Nec.	Childcare	Taxes

Notes: Observations are at the county level. Each column uses a different EPI Family Budget Calculator inflation measure (2023–2024, standardized). Panel A replicates Flavin's (2026) specification: OLS with controls for unemployment rate, bachelor's degree share, median age, evangelical congregation rate, percent white, and log population. The first SE row reports conventional OLS standard errors (as in the original); the second row reports heteroskedasticity-robust standard errors; the third row reports standard errors clustered at the state level. Panel B adds state fixed effects; robust standard errors in parentheses, cluster-bootstrapped standard errors (100 replications, clustered at the state level) in brackets. Stars on the coefficient correspond to the first SE row in each panel. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table J.2
Cross-Specification Robustness: Our Outcome with Flavin (2026) Controls

	(1) Total	(2) Housing	(3) Food	(4) Transport.	(5) Healthcare	(6) Other Nec.	(7) Childcare	(8) Taxes
Inflation	−0.0656***	−0.0163	0.0106	−0.0304*	−0.0422	−0.0470**	−0.0319**	−0.0318
	(0.0237)	(0.0155)	(0.0225)	(0.0168)	(0.0299)	(0.0236)	(0.0160)	(0.0257)
	[0.0353]	[0.0236]	[0.0324]	[0.0189]	[0.0346]	[0.0304]	[0.0249]	[0.0394]
Nominal Wage Growth	−0.0274*	−0.0269	−0.0273*	−0.0266	−0.0275*	−0.0289*	−0.0262	−0.0274*
	(0.0165)	(0.0165)	(0.0165)	(0.0164)	(0.0165)	(0.0165)	(0.0165)	(0.0165)
	[0.0152]	[0.0150]	[0.0147]	[0.0147]	[0.0150]	[0.0147]	[0.0150]	[0.0151]
Observations	3,098	3,098	3,098	3,098	3,098	3,098	3,098	3,098
R^2	0.404	0.403	0.402	0.403	0.403	0.403	0.403	0.403
State FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Flavin (2026)	Flavin (2026)	Flavin (2026)	Flavin (2026)	Flavin (2026)	Flavin (2026)	Flavin (2026)	Flavin (2026)
Wages	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Price measure	Total	Housing	Food	Transport.	Healthcare	Other Nec.	Childcare	Taxes

Notes: Observations are at the county level. Dependent variable is the change in two-party Republican presidential vote share, 2020–2024 (our measure). Controls are Flavin's (2026) set: unemployment rate (2023), bachelor's degree share, median age, evangelical congregation rate, percent white, and log population. Wage measure is change in average annual pay (QCEW), 2021–2024. Inflation is EPI Family Budget Calculator, 2021–2024. All variables standardized to zero mean and unit variance. Robust standard errors in parentheses; cluster-bootstrapped standard errors (100 replications, clustered at the state level) in brackets. All specifications include state fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

statistical significance relies on conventional OLS standard errors that ignore the spatial clustering inherent in MSA-level inflation measurement; once state-clustered standard errors are employed, most estimates become insignificant. The positive sign is driven by between-state covariation, and adding state fixed effects reverses it. Finally, omitting wages prevents distinguishing price salience from real income effects; when wages are included, the inflation coefficient is negative regardless of which control set is used.

References

Angrist, J.D., Pischke, J.S., 2009. Mostly Harmless Econometrics: An Empiricist's Companion. Princeton University Press, Princeton, NJ.

Autor, D., Dorn, D., Hanson, G., Majlesi, K., 2020. Importing political polarization? The electoral consequences of rising trade exposure. American Economic Review 110, 3139–3183.

Bureau of Economic Analysis, 2025. Methodology for regional price parities, real personal consumption expenditures, and real personal income. `https://www.bea.gov/sites/default/files/methodologies/Methodology-for-Regional-Price-Parities_0.pdf`. U.S. Department of Commerce, Bureau of Economic Analysis.

Charles, K.K., Hurst, E., Schwartz, M., 2019. The transformation of manufacturing and the decline of US employment. NBER Macroeconomics Annual 33, 307–372.

Dippel, C., Gold, R., Heblich, S., Pinto, R., 2022. The effect of trade on workers and voters. The Economic Journal 132, 199–217.

Economic Policy Institute, 2024. Family budget calculator: Frequently asked questions. https://www.epi.org/resources/budget/budget-faq/. Accessed April 2026.

Fiorina, M.P., 1981. Retrospective Voting in American National Elections. Yale University Press, New Haven, CT.

Flavin, P., 2026. Did Trump do better where inflation was worse? Evidence from county-level data. Electoral Studies 99, 103039.

Wolf, C.K., 2023. The missing intercept: A demand equivalence approach. American Economic Review 113, 2232–2269.