EconomiX

When Drought Travels: Partisan Electoral Responses Through Food Networks in U.S. Presidential Elections

Cécile Couharde
Rémi Generoso
Michaël Guillossou

2026-21 Document de Travail/ Working Paper

EconomiX - UMR 7235 Bâtiment Maurice Allais
Université Paris Nanterre 200, Avenue de la République
92001 Nanterre Cedex

Site Web : economix.fr
Contact : secreteriat@economix.fr

cnrs Université Paris Nanterre

When Drought Travels: Partisan Electoral Responses Through Food Networks in U.S. Presidential Elections

Cécile Couharde*[1], Rémi Generoso†[2], and Michaël Guillossou‡[1]

[1]EconomiX, Université Paris Nanterre

[2]LEM, Université de Lille; Chaire en macroéconomie et prévisions, Université du Québec Montréal (UQAM)

Abstract

This paper examines whether climate shocks are associated with electoral behavior beyond the counties in which they occur. The analysis covers 3,069 counties in the contiguous United States across five presidential elections from 2008 to 2024. County-level election returns are combined with pre-election drought measured using the Palmer Z Index and with network exposure constructed from estimated county-to-county food flows in 2007. A model in which voters interpret drought through pre-existing partisan preferences predicts that the electoral association of drought becomes more favorable to the Democratic candidate as a county's historical Democratic vote share increases. The estimates are consistent with this prediction. Local drought is associated with lower Democratic two-party vote shares in historically Republican counties and higher shares in historically Democratic counties. Network drought exhibits the same sign-changing pattern, including where measured local drought is absent. The estimated relationships cross zero at historical Democratic vote shares of about 0.40 for local drought and 0.44 for network exposure. Directional exposure through suppliers versus buyers adds little explanatory power conditional on total network exposure. The network relationship is comparatively stable across alternative specifications. Turnout and party-vote results suggest that the pattern reflects changes in party-specific electoral support rather than a common change in aggregate participation across counties.

Keywords: drought; food-flow networks; partisan heterogeneity; U.S. presidential elections

JEL: Q54; D72; Q13; L14; R15

*cecile.couharde@parisnanterre.fr

†remi.generoso@univ-lille.fr

‡michael.guillossou@parisnanterre.fr

1. Introduction

Climate shocks may affect political behavior beyond the locations in which they occur. Most studies of the electoral consequences of weather and natural disasters, however, define exposure locally: a county or region is considered affected when the physical event occurs within its boundaries (Achen and Bartels, 2004; Healy and Malhotra, 2009, 2010; Gasper and Reeves, 2011; Bechtel and Hainmueller, 2011; Cole et al., 2012; Hazlett and Mildenberger, 2020). A substantial part of this literature draws on a retrospective-voting framework in which voters reward or punish incumbent parties for locally disaster-related outcomes and government responses (Key, 1966; Kramer, 1971; Fair, 1978; Fiorina, 1981; Achen and Bartels, 2016).

This paper extends the analysis of electoral responses to climate shocks in two directions. First, it constructs a measure of drought exposure based on pre-sample food-flow connections and examines its relationship with county-level electoral outcomes separately from local drought. Second, it allows the relationships between local and network drought exposure and electoral outcomes to vary continuously with a county's historical Democratic vote share, rather than assuming a uniform electoral response to drought.

A theoretical framework links cross-county differences in these relationships to differences in how voters interpret drought-related information. Voters may draw different political conclusions from similar signals depending on their pre-existing partisan preferences (Bolsen et al., 2014; Bisgaard, 2015, 2019). Unlike conventional retrospective-voting models, which predict that adverse events reduce support for the incumbent party, the framework allows the association between drought and electoral outcomes to vary with a county's historical Democratic vote share. Its main empirical implication is that the relationship between drought and the Democratic vote share becomes more favorable as historical Democratic support increases. In the backlash case, this relationship is negative at low levels of historical Democratic support and positive at high levels, generating a sign reversal within the partisan distribution. If drought among food-flow partners is interpreted similarly to local drought, network exposure should display a similar pattern. These implications organize the empirical analysis.

We use a balanced panel of 3,069 counties in the contiguous United States covering five presidential elections between 2008 and 2024. We combine county-level election results with pre-election local drought, measured using the Palmer Z Index (Palmer, 1965). We also construct network drought exposure as a bilateral-flow-weighted average of drought among each county's food-flow partners, using 2007 county-to-county food-flow estimates from Karakoc et al. (2022), based on the Food Flow Model of Lin et al. (2019). We measure historical Democratic alignment using Democratic two-party vote shares from county-level Senate and gubernatorial elections held between 1992 and 2004. The regressions interact local and network drought exposure with this historical voting measure

to examine whether their associations with electoral outcomes vary across counties.

The estimates show that local and network drought exposure are associated with lower Democratic vote shares in historically Republican counties and higher shares in historically Democratic counties. In the baseline specification, a one-standard-deviation increase in local drought is associated with a 0.34-percentage-point decline in the Democratic vote share in counties where the historical Democratic vote share is 25 percent, and a 0.75-percentage-point increase where it is 75 percent. For a one-standard-deviation increase in mean network exposure, the corresponding estimates are approximately -0.41 and +0.68 percentage points. The estimated sign-change thresholds for local and mean network drought are approximately 0.40 and 0.44, respectively, close to the sample mean of historical Democratic support. The implied associations are therefore small near the center of the partisan distribution. The main empirical pattern is consequently one of partisan heterogeneity rather than a uniform electoral shift toward either party.

The network relationship is also present when measured local drought is absent. Among county-election observations with no local drought, high drought exposure among food-flow partners is associated with a negative difference in Democratic vote share at low levels of historical Democratic support and a positive difference at high levels. The relationship is also more pronounced at high than at low levels of network exposure. These results provide a more direct assessment of the network pattern because they separate drought among economically connected counties from contemporaneous drought within the county itself. We find little evidence that the association depends on whether drought is concentrated among net suppliers or net destination markets. The evidence is therefore stronger for overall drought exposure among connected counties than for differences based on the direction of food flows.

The network relationship is generally stable across alternative specifications. Allowing historical Democratic support to follow a differential trend in nonmetropolitan counties reduces the magnitude of the local-drought interaction but leaves the mean-network interaction largely unchanged. The main pattern also survives alternative election-specific fixed effects and checks for network concentration, spatial inference, and influential elections. By contrast, the local-drought interaction is more sensitive in magnitude and precision, particularly to the limited number of elections in the sample.

To assess whether the partisan pattern is consistent with simple retrospective punishment, we also compare drought relationships across elections held under Democratic and Republican administrations. Under a simple retrospective-voting interpretation, drought should be more unfavorable to the Democratic vote share during Democratic presidencies. The estimated differences do not have this predicted sign.

Supplementary outcomes help assess whether changes in the two-party vote share are associated with aggregate participation or with the partisan composition of votes. The estimated effects on total turnout relative to the citizen voting-age population are small and

statistically imprecise. The results for Democratic votes relative to this population show the same negative level coefficients and positive interactions as those for the Democratic two-party vote share. The corresponding results for Republican votes are weaker and less stable. Taken together, these estimates are consistent with changes in the partisan composition of votes rather than with a uniform change in aggregate turnout.

This paper contributes to three strands of literature. First, it extends research on the electoral effects of weather shocks and natural disasters by considering exposure beyond the location of the physical event. Existing studies generally focus on voters in directly affected areas. This paper shows that drought among economically connected counties is associated with electoral outcomes even in county-election observations with no measured local drought. The political relevance of climate shocks may therefore extend beyond directly affected areas.

Second, the paper contributes to research on how shocks propagate through production networks. Existing studies show that shocks can spread through input–output linkages and affect production, employment, prices, and aggregate activity (Acemoglu et al., 2012; Barrot and Sauvagnat, 2016; Boehm et al., 2019; Carvalho et al., 2021). Related work shows that geographically concentrated trade shocks can also influence electoral behavior (Autor et al., 2020). By linking drought among food-flow partners to electoral outcomes in connected counties, including counties without measured local drought, this paper brings a network perspective to the political consequences of climate shocks.

The third contribution relates to partisan differences in responses to climate information and adverse events. Existing studies show that partisans differ in their beliefs about climate change, their processing of climate-related information, and their responses to environmental shocks (Kahan et al., 2012; Bolsen et al., 2014; Bisgaard, 2015, 2019; Tesler, 2018; Druckman and McGrath, 2019; Merkley and Stecula, 2018, 2021). Related evidence indicates that direct exposure to adverse environmental events can produce different electoral responses in Republican- and Democratic-leaning areas (Hazlett and Mildenberger, 2020). Rather than dividing counties into only broad Republican and Democratic groups, we allow the relationship between drought exposure and the Democratic vote share to vary across the full distribution of counties' historical Democratic vote shares. The estimates show a gradual change from a negative relationship at low historical Democratic vote shares to a positive relationship at high shares. This pattern appears for both local drought and drought among food-flow partners and differs from the prediction of a simple incumbent-punishment model.

The remainder of the paper is organized as follows. Section 2 presents the partisan-attribution framework and derives its empirical implications. Section 3 describes the electoral, drought, historical-partisanship, and food-flow data. Section 4 presents the empirical specification and identification strategy, including the augmented specification that allows for differential partisan trends in nonmetropolitan counties. Section 5 reports

the baseline estimates, examines heterogeneity by historical Democratic support, and assesses the network relationship among observations with no measured local drought. Section 6 examines the robustness and sensitivity of the main estimates to alternative election-specific fixed effects, network characteristics, spatial inference, and election samples. Section 7 examines incumbent-party heterogeneity and the electoral channels underlying the two-party vote-share results. Section 8 concludes.

2. A Model of Partisan Interpretation of Local and Network Drought

This section presents a model in which prior partisan preferences shape how voters interpret drought and derives the vote-share equation used in the empirical analysis. The main text presents the model and the restrictions tested in the empirical analysis; the formal derivations and proofs are provided in Appendix A.

2.1 Retrospective punishment versus partisan attribution

The mechanism developed here differs from a standard retrospective-voting framework. In such models, adverse outcomes provide information about the incumbent party's performance, leading voters to reward or punish the party in office (Key, 1966; Kramer, 1971; Fair, 1978; Fiorina, 1981). Applied to weather, this logic implies that voters may punish the incumbent party for drought, even when the government did not cause it (Achen and Bartels, 2016).

The framework developed here treats drought differently. Its political implications are not fixed but depend on voters' prior beliefs. Voters observe the same weather event and agree that it occurred, but they differ in whether they regard it as informative about climate risks that public policy can address. The disagreement is therefore not about the drought itself, but about what it signals about climate risk.

These differences arise because partisans may interpret the same drought evidence differently when updating their beliefs. Beliefs about climate risk and support for policies to address it are closely related to partisanship (Kahan et al., 2012; Merkley and Stecula, 2018; Tesler, 2018). Partisan groups may therefore revise their climate-related beliefs–and the importance they place on climate policy–by different amounts or even in opposite directions (Bolsen et al., 2014; Bisgaard, 2015, 2019; Bullock et al., 2015; Druckman and McGrath, 2019; Bonomi et al., 2021). Evidence that local weather conditions influence Americans' views about global warming supports the idea that an external weather event can change climate-related beliefs (Egan and Mullin, 2012; Deryugina, 2013). The model therefore does not assume either a uniform electoral response or a fixed direction of the relationship; its magnitude and sign are estimated empirically.

2.2 Partisan belief updating

Consider a voter with partisan identity $\sigma \in D, R$ who observes a drought of intensity $D \geq 0$ and assigns probability $\mu_\sigma(D)$ to the drought arising from a policy-relevant climate state M, rather than a natural state O. State M means only that the drought is viewed as informative about climate-related risks that public policy–such as adaptation, relief, or mitigation–may address. The voter need not conclude that anthropogenic climate change caused this particular drought or that the incumbent government is responsible for it.

Let $\ell(D) = f_M(D)/f_O(D)$ denote the likelihood ratio of observing drought intensity D under state M rather than state O. Because $\ell'(D) > 0$, more severe droughts provide stronger evidence in favor of state M.

Let $\Pr(M)$ and $\Pr(O)$ denote the common prior probabilities of states M and O, respectively, with $\Pr(M) + \Pr(O) = 1$. The common prior odds of state M relative to state O are then given by $\omega_0 = \Pr(M)/\Pr(O)$.

A voter's posterior belief after observing the drought combines the common prior odds, the evidence conveyed by the drought, and an adjustment linked to partisan identity:

$$\ln \frac{\mu_\sigma(D)}{1 - \mu_\sigma(D)} = \ln \omega_0 + \ln \ell(D) + m_\sigma(D) \tag{1}$$

The first two terms represent standard Bayesian updating: voters combine the common prior odds with the evidence provided by the drought. The third term captures the adjustment associated with partisan identity.

This adjustment is specified as $m_D(D) = \alpha_D + \gamma_D D$ for Democratic voters and $m_R(D) = -\alpha_R - \gamma_R D$ for Republican voters, where $\alpha_D, \alpha_R, \gamma_D, \gamma_R \geq 0$. The parameters α_D and α_R capture partisan differences in the tendency to view climate information as relevant to policy, independent of drought intensity. The parameters γ_D and γ_R capture how these tendencies strengthen as drought intensity increases. This specification reflects evidence that partisans may agree on observed conditions but draw different explanations and policy implications, and that political identity can shape the interpretation of new information in ways consistent with prior group beliefs (Bisgaard, 2015, 2019; Bolsen et al., 2014; Bonomi et al., 2021).

Under the imposed sign restrictions, greater drought intensity leads Democratic voters to place more weight on the drought as evidence of policy-relevant climate risk, while Republican voters place less weight on this interpretation. This does not necessarily imply that Republicans become less likely to view climate risk as policy-relevant as drought severity increases. The direct informational effect of drought, captured by $\ln \ell(D)$, also increases with drought intensity and may outweigh the negative partisan identity component. The specification therefore allows partisanship to amplify belief updating among Democrats and dampen–or, if sufficiently strong, reverse–it among Republicans.

To link beliefs to voting behavior, we assume that the Democratic platform becomes more attractive relative to the Republican platform when voters view drought as evidence of policy-relevant climate risk:

$$\kappa \equiv [V_D(M) - V_R(M)] - [V_D(O) - V_R(O)] > 0 \tag{2}$$

where κ measures the increase in the Democratic platform's relative advantage under state M compared with state O.

Thus, the condition $\kappa > 0$ does not require every voter to prefer the Democratic platform. It requires only that a higher belief in state M makes the Democratic platform relatively more attractive. This assumption is consistent with evidence that Democratic voters and elites are generally more supportive of climate policy and more likely than Republicans to view climate risks as a policy priority (Egan and Mullin, 2017).

For a voter with partisan identity σ, the resulting change in the relative attractiveness of the Democratic platform is:

$$A_\sigma(D) = \kappa \left[\mu_\sigma(D) - \mu_\sigma(0)\right] \tag{3}$$

The term $\mu_\sigma(D) - \mu_\sigma(0)$ measures how drought intensity D changes the voter's belief that the event reflects a policy-relevant climate risk. The parameter κ determines how strongly this belief change increases the attractiveness of the Democratic platform relative to the Republican platform.

Differentiating this expression gives (Lemma A.1),

$$sign\left(\frac{dA_\sigma}{dD}\right) = sign(s(D) + m'_\sigma(D)) \tag{4}$$

with $s(D) \equiv \frac{d \ln \ell(D)}{dD} > 0$.

The term $s(D)$ measures the common informational effect of greater drought intensity: because $s(D) > 0$, a more severe drought provides stronger evidence in favor of the policy-relevant climate state M. The term $m'_\sigma(D)$ captures how this response varies with partisan identity. Thus, the effect of drought intensity on the relative attractiveness of the Democratic platform depends on whether the informational and partisan components reinforce or offset each other.

At a reference drought intensity $\bar{D}$, define

$$\delta_D \equiv \left.\frac{dA_D(D)}{dD}\right|_{D=\bar{D}} > 0 \qquad \text{and} \qquad \delta_R \equiv -\left.\frac{dA_R(D)}{dD}\right|_{D=\bar{D}}. \tag{5}$$

The parameter δ_D (δ_R) measures the marginal increase (decline) in the relative attractiveness of the Democratic platform among Democratic- (Republican-) identified voters as

drought intensity increases, evaluated at $\bar{D}$.

Because $m'_D = \gamma_D$, the common informational effect and the Democratic identity adjustment reinforce one another. A more severe drought therefore increases the relative attractiveness of the Democratic platform among Democratic-identified voters.

By contrast, because $m'_R = -\gamma_R$, the Republican identity adjustment offsets the common informational effect. Whether a more severe drought increases or decreases the relative attractiveness of the Democratic platform among Republican-identified voters therefore depends on the sign of $s(\bar{D}) - \gamma_R$:

- When $\gamma_R < s(\bar{D})$, $\delta_R < 0$: greater drought intensity increases the relative attractiveness of the Democratic platform among Republican-identified voters because the informational effect outweighs the opposing partisan component.
- When $\gamma_R = s(\bar{D})$, $\delta_R = 0$: the two components offset each other, so drought intensity does not change the relative attractiveness of the Democratic platform among Republican-identified voters.
- When $\gamma_R > s(\bar{D})$, $\delta_R > 0$: the partisan component dominates the informational effect, so greater drought intensity reduces the relative attractiveness of the Democratic platform among Republican-identified voters. This case captures a partisan backlash, as documented in the climate-communication and disaster-exposure literature (Hart and Nisbet, 2012; Merkley and Stecula, 2021; Djourelova et al., 2023).

Among Democratic-identified voters, greater drought intensity unambiguously increases the relative attractiveness of the Democratic platform because the partisan and informational components reinforce each other. Among Republican-identified voters, the direction of this response is ambiguous because the two components work in opposite directions. County-level effects therefore depend on the relative shares of Democratic- and Republican-identified voters.

2.3 Aggregation to the county level

The effect of drought on the county-level Democratic vote share is a weighted average of the changes in support for the Democratic candidate within the Democratic- and Republican-identified groups, with each group weighted by its share of the county electorate. Let π_i denote the historical Democratic vote share in county i. We use π_i as a proxy for the share of Democratic-identified voters in county i, s_{Di}, so that $s_{Di} \approx \pi_i$. The share of Republican-identified voters in county i, s_{Ri}, is therefore approximated by $s_{Ri} \approx 1 - \pi_i$.[1]

[1]We use historical Democratic vote share as a proxy for the share of Democratic-identified voters rather than as a direct measure of partisan identification. Historical vote shares may also reflect differences in turnout, support among independent voters, candidate-specific factors, and the treatment of uncontested elections. The empirical results should therefore be interpreted as heterogeneity across counties with different levels of historical Democratic alignment, rather than as direct estimates of responses among Democratic- and Republican-identified voters.

In a probabilistic-voting framework, voter v with partisan identity σ votes for the Democratic candidate with probability $\Lambda\left(u_v + A_\sigma(D)\right)$, where $A_\sigma(D)$ captures how drought changes the relative attractiveness of the Democratic platform for that partisan group and u_v the voter's baseline support for the Democratic candidate. Because $\Lambda(\cdot)$ is increasing, the probability of voting for the Democratic candidate rises with the voter's underlying preference for the Democratic candidate, u_v, and with any increase in the Democratic platform's relative attractiveness due to drought, $A_\sigma(D)$.

Averaging the marginal effect of drought across voters and incorporating how changes in the relative attractiveness of the Democratic platform affect voting probabilities into the group-specific parameters yields the county-level effect of drought intensity, D_{it}, on the Democratic vote share, y_{it} (Proposition A.2):

$$\frac{\partial y_{it}}{\partial D_{it}} = \pi_i \delta_D - (1 - \pi_i)\,\delta_R = -\delta_R + (\delta_D + \delta_R)\,\pi_i = \beta_0 + \beta_1 \pi_i \tag{6}$$

The predicted change in the county's Democratic vote share as drought intensity increases is therefore linear in its historical Democratic vote share.[2] The coefficient $\beta_0 = -\delta_R$ gives this predicted change when $\pi_i = 0$, corresponding to a hypothetical county with no historical Democratic support. The coefficient $\beta_1 = \delta_D + \delta_R$ measures how the predicted change varies as the county's historical Democratic vote share increases.

Consistent with evidence that climate-related events and information tend to produce stronger pro-climate responses among Democrats than among Republicans, we focus on the case in which $-\delta_R \leq \delta_D$ (Hazlett and Mildenberger, 2020; Merkley and Stecula, 2021). This restriction rules out a stronger pro-Democratic response among Republican-identified voters than among Democratic-identified voters.

Under this restriction, three cases can be distinguished.

- Under a uniform pro-Democratic response, drought shifts Democratic and Republican voters equally toward the Democratic candidate. This requires $-\delta_R = \delta_D > 0$, so that $\beta_0 = -\delta_R > 0$ and $\beta_1 = 0$. The predicted increase in the Democratic vote share is then the same across counties, regardless of their partisan composition.
- Under a weaker pro-Democratic response among Republican voters, Republican voters also shift toward the Democratic candidate, but by less than Democratic voters. This requires $0 \leq -\delta_R < \delta_D$, so that $\beta_0 = -\delta_R \geq 0$ and $\beta_1 = \delta_D + \delta_R > 0$. The predicted change in the Democratic vote share is therefore positive across counties but smaller in more Republican counties.

[2]The linear relationship in π_i follows from aggregating the group-specific responses using historical Democratic support as a proxy for the relative shares of Democratic- and Republican-identified voters. The linear form itself is therefore not specific to the belief-updating mechanism. The more informative implication of the framework concerns the sign and ordering of the responses. In the backlash case, the relationship between drought and the Democratic vote share is negative at low levels of historical Democratic support and becomes more positive as historical Democratic support increases, eventually changing sign within the partisan distribution.

- Under a backlash among Republican voters, drought reduces their support for the Democratic candidate. This requires $\delta_R > 0$, implying that $\beta_0 = -\delta_R < 0$ and $\beta_1 = \delta_D + \delta_R > 0$. The predicted change is therefore negative in counties with a sufficiently low historical Democratic vote share and becomes more favorable as the historical Democratic vote share increases.

The predicted effect of greater drought intensity on the Democratic vote share is therefore negative when the Democratic share is sufficiently low and becomes positive only after crossing the threshold:

$$\pi_D^* = -\frac{\beta_0}{\beta_1} = \frac{\delta_R}{\delta_D + \delta_R} \in [0, 1] \tag{7}$$

The threshold π_D^* is the historical Democratic vote share at which the positive response among Democratic-identified voters exactly offsets the negative response among Republican-identified voters. When $\pi_i < \pi_D^*$, the Republican backlash dominates, and greater drought intensity lowers the county's Democratic vote share. When $\pi_i > \pi_D^*$, the response among Democratic-identified voters dominates, and greater drought intensity raises the Democratic vote share. At $\pi_i = \pi_D^*$, the predicted change is zero.

These alternatives can be distinguished by the estimated relationship between drought and the Democratic vote share and how it varies with counties' historical partisan composition. The theoretical framework therefore predicts that the relationship should become more favorable to Democrats as the historical Democratic vote share increases. In the backlash case, where greater drought intensity reduces Democratic support among Republican-identified voters, the relationship should be negative in counties whose historical Democratic vote share falls below the threshold π_D^*.

2.4 Transmission through food-flow networks

Drought may affect a county not only directly but also through economic links to drought-affected counties, consistent with research showing that shocks propagate through production networks (Acemoglu et al., 2012; Barrot and Sauvagnat, 2016; Boehm et al., 2019; Carvalho et al., 2021). Using food-flow weights, we define county i's network exposure as the weighted average of drought intensity among its food-flow partners, with greater weight assigned to more strongly connected counties:

$$N_{it}^m = \sum_{j \neq i} w_{ij}^m D_{jt} \tag{8}$$

D_{jt} is drought intensity in partner county j, and w_{ij}^m the strength of the food-flow link between counties i and j.

Economic transmission does not necessarily lead to political belief updating. Drought

in a partner county may affect local prices, income, or employment without being viewed by voters as evidence of climate risk. We therefore define the climate-related signal perceived in county i as:

$$\tilde{D}_{it}^{m} = D_{it} + \varphi N_{it}^{m} \tag{9}$$

where the parameter $\varphi > 0$ measures how strongly network drought exposure contributes to the perceived climate signal relative to local drought. It therefore captures the climate relevance that voters attach to network exposure. The parameter is assumed to be the same for Democratic and Republican voters.

When local and network drought exposure operate through the same belief-updating mechanism, and network exposure is scaled by the common parameter φ, the effect of network drought exposure on the Democratic vote share is given by (Proposition A.3):

$$\frac{\partial y_{it}}{\partial N_{it}^{m}} = \varphi \left(\beta_0 + \beta_1 \pi_i\right) \tag{10}$$

Thus, the network effect varies with the county's historical Democratic vote share in the same way as the local effect, but its magnitude is scaled by φ. When $\varphi > 0$, the effect of network drought exposure on the Democratic vote share changes sign at the following threshold:

$$\pi_{N^m}^{*} = -\frac{\varphi\beta_0}{\varphi\beta_1} = -\frac{\beta_0}{\beta_1} = \pi_{D}^{*} \tag{11}$$

The model therefore predicts that local and network drought exposure have a positive relationship with the Democratic vote share above the same partisan threshold.

If network drought exposure is weighted differently by Democratic- and Republican-identified voters, its effect on the county-level Democratic vote share is:

$$\frac{\partial y_{it}}{\partial N_{it}^{m}} = \pi_i \varphi_D \delta_D - \left(1 - \pi_i\right) \varphi_R \delta_R \tag{12}$$

Setting the effect of network drought exposure on the Democratic vote share equal to zero yields the following threshold:

$$\pi_{N^m}^{*} = \frac{\varphi_R \delta_R}{\varphi_D \delta_D + \varphi_R \delta_R} \tag{13}$$

In this case, the effects of local and network drought exposure need not become zero at the same level of historical Democratic support. A difference between the two thresholds would therefore challenge the assumption that network exposure receives the same weight across partisan groups. It would not, however, reject the broader belief-updating mechanism through which perceived drought exposure shapes voters' political choices.

Finally, the belief-updating mechanism is assumed to depend on total drought exposure

across a county's economic network, regardless of whether the affected partners are suppliers to county i or buyers from it. Under this assumption, once network exposure N_{it}^{m} is included, the difference between exposure through suppliers and buyers, N_{it}^{Δ}, should have no additional effect on the Democratic vote share. A statistically significant effect of N_{it}^{Δ} would instead suggest that the direction of transmission matters. For example, drought among suppliers may raise input prices or disrupt production, whereas drought among buyers may reduce demand for local output. Such a result would indicate an additional channel linked to the direction of food flows, rather than reject the broader idea that network drought exposure shapes political beliefs.

2.5 Empirical implications and testable restrictions

Linearizing individual voting responses around the reference drought intensity $\bar{D}$ and adding other determinants of voting behavior yields the following empirical specification:

$$y_{it} = \alpha_i + \lambda_t + (\beta_0 + \beta_1 \pi_i) D_{it} + \varphi (\beta_0 + \beta_1 \pi_i) N_{it}^{m} + u_{it} \tag{14}$$

where y_{it} is the Democratic vote share in county i and election t, while α_i and λ_t are county and election fixed effects, respectively. County fixed effects absorb persistent differences across counties, while election fixed effects capture factors common to all counties in each election.

Equation (14) shows that local drought D_{it} and network drought exposure N_{it}^{m} share the same term, $\beta_0 + \beta_1 \pi_i$. Their effects on the Democratic vote share therefore vary with the county's historical Democratic vote share, π_i, with the magnitude of network drought exposure scaled by φ.

For estimation, Section 4 embeds Equation (14) in a more flexible specification that allows local drought, total network exposure, and directional network exposure to have separate main and interaction coefficients. For the three channels, $k \in \{D_{it}, N_{it}^{m}, N_{it}^{\Delta}\}$, the coefficients β_{0D}, β_{0N}^{m}, and β_{0N}^{Δ} measure the estimated relationship when the county's historical Democratic vote share is zero. The coefficients β_{1D}, β_{1N}^{m}, and β_{1N}^{Δ} measure how each relationship changes as the historical Democratic vote share increases.

The theoretical framework implies three testable restrictions:

- First, the relationships between local drought or network drought exposure and the Democratic vote share should become more positive as historical Democratic support increases: $\beta_{1D} > 0$ and $\beta_{1N}^{m} > 0$. In the backlash case, the predicted relationships in counties with no historical Democratic support should be negative: $\beta_{0D} < 0$ and $\beta_{0N}^{m} < 0$ when $\pi_i = 0$.
- Second, if network exposure is scaled proportionally and receives the same weight across partisan groups, local drought and network drought exposure should become

positive at the same level of historical Democratic support: $\beta^m_{0N}/\beta^m_{1N} = \beta_{0D}/\beta_{1D}$. This restriction can equivalently be tested as $\beta^m_{0N}\beta_{1D} - \beta^m_{1N}\beta_{0D} = 0$. When the relevant coefficients are nonzero, the common scaling factor is $\varphi = \beta^m_{0N}/\beta_{0D} = \beta^m_{1N}/\beta_{1D}$.

- Third, conditional on total network exposure, N^m_{it}, the supplier–buyer composition of that exposure should have no additional relationship with the Democratic vote share: $\beta^{\Delta}_{0N} = \beta^{\Delta}_{1N} = 0$.

The first restriction characterizes the partisan-gradient pattern implied by the framework.[3] In the backlash case, the more distinctive empirical implication is the combination of a positive partisan gradient and a sign reversal within the observed distribution of historical Democratic support. The second tests the stronger prediction that local and network drought operate through the same mechanism and differ only by a common scaling factor. The third examines whether exposure through suppliers differs from exposure through buyers. Section 4 presents the unrestricted estimating equation and explains how these restrictions are tested.

3. Data and Measurement

3.1 Sample and electoral outcomes

The empirical panel contains 3,069 counties in the contiguous United States observed across five presidential elections from 2008 through 2024, yielding 15,345 county-election observations. The sample is restricted to counties covered by the Census of Agriculture conducted by the United States Department of Agriculture (USDA) because the county-to-county food-flow estimates are based on recorded agricultural activity. This restriction excludes Virginia's independent cities and a small number of highly urbanized counties with no recorded agricultural production.

The primary dependent variable is the Democratic two-party vote share in presidential elections. County-level returns through 2020 are drawn from Amlani and Algara (2021), while 2024 returns are obtained from the MIT Election Data and Science Lab.[4] To clarify whether changes in the two-party vote share reflect turnout or the partisan composition of votes, we also examine turnout, Democratic votes, and Republican votes, each as a share of the citizen voting-age population.

[3] Other mechanisms that vary systematically with historical partisan alignment—including heterogeneous economic exposure, partisan mobilization, or differences in media and elite cues—could generate similar reduced-form patterns. The model should therefore be interpreted as an organizing framework for the empirical heterogeneity rather than as a uniquely identified microfoundation.

[4] See https://electionlab.mit.edu/data.

3.2 Local drought exposure

Drought is measured using the monthly county-level Palmer Z Index (Palmer, 1965) from the National Oceanic and Atmospheric Administration's National Centers for Environmental Information (NOAA/NCEI) nClimDiv dataset.[5] The Palmer Z Index measures monthly moisture conditions relative to local historical norms, with negative values indicating moisture deficits.

To construct an election-specific measure, let Z_{it} denote the day-weighted average Palmer Z Index for county i over the six months ending in September of election year t. We define the nonnegative drought measure as $\widetilde{D}_{it} = \max\{-Z_{it}, 0\}$. Thus, $\widetilde{D}_{it}$ equals zero when the six-month average shows no moisture deficit and rises with deficit intensity.

We then standardize the measure by its standard deviation in the estimation sample $D_{it} = \widetilde{D}_{it}/s_{\widetilde{D}}$. The estimated coefficients on the drought variables can therefore be interpreted as the estimated change in the Democratic vote share associated with a one-standard-deviation increase in pre-election drought intensity.

The six-month window covers much of the growing and harvest period and ends in September, before the November election. The measure therefore uses only climatic conditions observed before voting takes place. Drought during this period can reduce crop yields, making its agricultural consequences more visible before the election (Kuwayama et al., 2019).

Table 1 describes the incidence and intensity of pre-election drought across presidential election years. The 2012 election stands out, with 81.2 percent of counties experiencing positive drought exposure and mean standardized intensity reaching 1.529. This exceptional exposure coincides with the severe and rapidly developing 2012 Central U.S. flash drought (Hoerling et al., 2014; Otkin et al., 2016; Basara et al., 2019). Given the unusually widespread and severe drought conditions in 2012, Section 6.3 examines whether the benchmark estimates are disproportionately influenced by this election by re-estimating the model after excluding 2012.

[5] See https://www.ncei.noaa.gov/access/metadata/landing-page/bin/iso?id=gov.noaa.ncdc:C00005.

Table 1: Pre-election drought incidence and intensity, by presidential election year

Election year	Share with drought > 0	Standardized drought		N
		Mean	S.D.	
2008	0.464	0.440	0.663	3,069
2012	0.812	1.529	1.260	3,069
2016	0.437	0.564	0.836	3,069
2020	0.400	0.549	0.942	3,069
2024	0.476	0.503	0.744	3,069
Full sample	0.518	0.717	1.000	15,345

Note: The table reports drought incidence and standardized intensity by election year. Drought is constructed from the day-weighted Palmer Z Index over the six months ending in September, with larger values indicating more severe moisture deficits, and is standardized over the estimation sample.

3.3 Food-flow network exposure

Network variables are constructed from county-to-county food-flow estimates developed by Karakoc et al. (2022). These data allocate Freight Analysis Framework flows to county pairs using the Food Flow Model of Lin et al. (2019). We use flows from 2007, before the first election in our sample, and aggregate SCTG categories 02–04: cereal grains (including seed) (SCTG 02); other agricultural products, excluding animal feed, cereal grains, and forage products (SCTG 03); and animal feed, eggs, honey, and other products of animal origin (SCTG 04). These categories consist mainly of primary or minimally processed agricultural commodities whose production is directly exposed to growing-season weather.

Appendix Table B.1 reports the composition of the resulting network. Cereal grains account for 71.2 percent of the selected 2007 flows, while SCTG 03 and SCTG 04 account for 17.0 and 11.8 percent, respectively. Thus, the selected network is predominantly composed of commodity flows whose origin is closely linked to primary agricultural production.[6]

Let F_{ij}^{g} denote the 2007 flow of commodity category g from county i to county j. We

[6]We exclude SCTG 01—live animals and fish—because shipments in this category reflect heterogeneous biological production cycles, inventory adjustments, and herd-management decisions rather than a contemporaneous growing-season production response comparable to that of crops. Biological lags are particularly important in livestock supply (Chavas, 2000; Aadland, 2004). We also exclude SCTG 05–07—meat, poultry, fish, and seafood preparations; milled-grain and bakery products; and other prepared foodstuffs, fats, and oils—because these categories comprise products that have undergone downstream transformation. In the Food Flow Model, commodity demand includes intermediate uses such as food processing and feed use (Lin et al., 2019). The origin county of a processed shipment may therefore identify the location of a slaughterhouse, mill, refinery, or other processing establishment rather than the location where its agricultural inputs were produced. This issue is especially relevant for products that combine inputs from several locations and pass through multiple stages of the supply chain (Bruckner et al., 2019). Restricting the network to SCTG 02–04 therefore improves the link between origin-county weather and primary agricultural production.

aggregate flows across the selected SCTG categories as

$$F_{ij} = \sum_{g \in \{02,\,03,\,04\}} F_{ij}^{g}, \tag{15}$$

and set all within-county flows, F_{ii}, equal to zero.

The total bilateral flow intensity between county i and partner j is $B_{ij} = F_{ij} + F_{ji}$. Let $S_i = \sum_{j \neq i} B_{ij}$ denote county i's total bilateral flow intensity across all partner counties. For counties with $S_i > 0$, the normalized bilateral network weight assigned to partner county j is

$$w_{ij}^{m} = \frac{B_{ij}}{S_i}. \tag{16}$$

For the 13 counties with no external SCTG 02–04 bilateral flows ($S_i = 0$), we set all network weights equal to zero. Their network drought exposure is consequently also equal to zero.

Table 2 describes the structure of the resulting pre-sample network.

Table 2: Characteristics of the 2007 food-flow network

Variable	**Mean**	**SD**	**P10**	**Median**	**P90**
Number of connected partner counties	53.847	45.837	9.000	44.000	117.200
Total bilateral food flow (10^9 kg)	1.091	1.622	0.022	0.543	2.817
Largest-partner share	0.437	0.257	0.151	0.374	0.865
Five-largest-partner share	0.798	0.205	0.480	0.862	1.000
Herfindahl index of partner weights	0.318	0.260	0.069	0.229	0.740
Flow-weighted partner distance (km)	360.3	317.1	126.6	265.0	701.3

Notes: Statistics are reported across the 3,069 counties in the baseline sample. Bilateral flow is $B_{ij} = F_{ij} + F_{ji}$ and is based on 2007 SCTG 02–04 flows, excluding within-county flows. Partner shares are expressed relative to total external bilateral flow, and partner distance is bilateral-flow weighted. For the 13 counties with no external flows, network weights and network exposure are set to zero. Total flow is reported in billions of kilograms.

The typical county is connected to a substantial number of other counties: the median number of external partners is 44, and the mean is 53.8. At the same time, bilateral flows are relatively concentrated. For the median county, the largest partner accounts for 37.4 percent of total bilateral flow, while the five largest partners together account for 86.2 percent. The network also extends beyond immediately neighboring counties: the median flow-weighted partner distance is 265 km, and the 90th percentile is approximately 701 km. These features imply that network exposure aggregates drought conditions across multiple connected locations while assigning substantial weight to a relatively small set of economically important partners. Because this concentration raises the possibility that the estimates could be disproportionately influenced by network hubs or counties with unusually concentrated partner weights, we assess the robustness of the results to excluding such counties in Section 6.2.

The bilateral-flow-weighted mean drought exposure of county i's network is then defined as

$$N_{it}^{m} = \sum_{j \neq i} w_{ij}^{m} D_{jt}, \tag{17}$$

where D_{jt} denotes drought intensity in partner county j at time t. For counties with positive external flows, N_{it}^{m} is therefore a weighted average of drought intensity across food-flow partners, with weights summing to one.

The second network variable captures the direction of bilateral food flows while retaining the same normalization factor, S_i. For $S_i > 0$, the directional asymmetry weight is defined as

$$w_{ij}^{\Delta} = \frac{F_{ji} - F_{ij}}{S_i}. \tag{18}$$

As with the mean-network weights, w_{ij}^{Δ} is set to zero when $S_i = 0$.

A positive value means that county j is a net supplier to county i, whereas a negative value means that county j is a net destination for flows originating in county i.

The asymmetric network exposure is

$$N_{it}^{\Delta} = \sum_{j \neq i} w_{ij}^{\Delta} D_{jt}. \tag{19}$$

Higher values of N_{it}^{Δ} indicate that drought exposure is more concentrated among net suppliers than among net destination markets.

The distinction between N_{it}^{m} and N_{it}^{Δ} allows the empirical analysis to separate overall drought exposure among food-flow partners from differences in exposure between suppliers and destination markets. A mechanism operating specifically through upstream suppliers or downstream destination markets should therefore generate a relationship between N_{it}^{Δ} and Democratic vote share. By contrast, broader economic, informational, or spatial spillovers that depend mainly on the strength of food-flow links may operate through N_{it}^{m}.

3.4 Historical Democratic support and regional drought

Historical Democratic support, π_i, is constructed using county-level results from Senate and gubernatorial elections held between 1992 and 2004 (Amlani and Algara, 2021), before the presidential elections in the estimation sample.

For each office type $o \in \{sen, gov\}$, we calculate the average Democratic share of the two-party vote across all elections of that type observed in county i:

$$\pi_i^{o} = \frac{1}{|T_i^{o}|} \sum_{e \in T_i^{o}} \frac{V_{ie}^{D}}{V_{ie}^{D} + V_{ie}^{R}}, \qquad 1992 \leq year(e) \leq 2004, \tag{20}$$

where T_i^{o} denotes the set of elections for office type o observed in county i, and V_{ie}^{D} and

V_{ie}^R denote the votes received by the Democratic and Republican candidates, respectively. Higher values of π_i^o indicate stronger historical Democratic support.

We then construct the overall measure by assigning equal weight to the Senate and gubernatorial components:

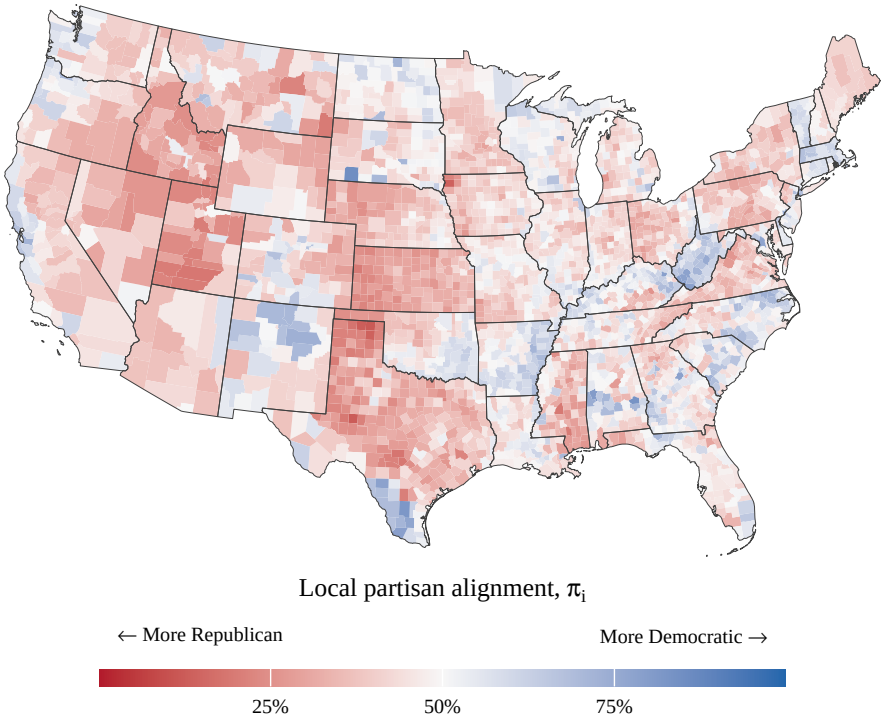
$$\pi_i = \tfrac{1}{2}\left(\pi_i^{sen} + \pi_i^{gov}\right). \tag{21}$$

We retain uncontested elections because excluding them would disproportionately remove counties with strong partisan support and shift the measure toward more electorally competitive places. When one major party does not field a candidate, its vote total is recorded as zero. Such cases are informative about the party's local electoral strength.

The evidence suggests that historical Democratic support is not simply a proxy for rurality. The distributions of π_i overlap substantially across metropolitan and nonmetropolitan counties. The difference in mean support is only 0.014, corresponding to a standardized mean difference of 0.13 (Appendix Table B.2). Moreover, a regression of π_i on a nonmetropolitan indicator yields an R^2 of just 0.36%, indicating that metropolitan–nonmetropolitan status explains only a negligible share of the cross-county variation in historical Democratic support.

Figure 1 shows a similar pattern: counties with higher and lower historical Democratic support are found in both metropolitan and nonmetropolitan areas. The measure therefore does not simply reproduce the rural–urban classification.

Figure 1: Historical local partisan alignment before 2008, by county, 1992–2004



Note: The map reports π_i, constructed from county-level Democratic two-party vote shares in statewide Senate and gubernatorial elections held between 1992 and 2004. Higher values indicate stronger historical Democratic alignment.

Because drought conditions are spatially correlated, the empirical specification includes a control for the county's broader regional drought environment. For each county i, D_{it}^{reg} is defined as the unweighted average drought intensity among all other counties located within 1,000 km of county i:

$$D_{it}^{reg} = \frac{1}{|R_i|} \sum_{j \in R_i} D_{jt}, \qquad R_i = \{j \neq i : d_{ij} \leq 1000 \text{ km}\}, \tag{22}$$

where d_{ij} denotes the Euclidean distance between the medoids of counties i and j. County i is excluded from the reference set R_i, and every county within the 1,000 km radius receives equal weight; the measure is therefore neither population-weighted nor distance-weighted.

The 1,000 km radius is deliberately broad. Its purpose is to capture spatially correlated drought conditions shared across both nearby and more distant counties. Conditional on D_{it}^{reg}, the coefficients on N_{it}^{m} and N_{it}^{Δ} are estimated from variation in food-flow-weighted drought intensity that remains after accounting for average drought conditions around county i. Thus, if the relationship between the outcome and N_{it}^{m} persists after controlling for D_{it}^{reg}, it is less likely to reflect solely the geographic proximity of economically connected counties and their resulting exposure to correlated drought conditions.

Appendix Table B.3 reports descriptive statistics for the variables entering the baseline

specification, including local drought, mean and directional network exposure, and regional drought.

4. Empirical Strategy and Identification

The baseline specification estimates the relationships between local and food-network drought and the Democratic two-party vote share, allowing each relationship to vary with the county's historical Democratic vote share:

$$y_{it} = \alpha_i + \lambda_t + \sum_k X_{it}^k \left[\beta_{0k} + \beta_{1k}\pi_i\right] + \theta D_{it}^{reg} + \eta\, \pi_i\Big(t - \bar{t}\Big) + \varepsilon_{it}. \tag{23}$$

where X_{it}^k belongs to $\{D_{it},\, N_{it}^m,\, N_{it}^{\Delta}\}$. Specifically, D_{it} is local pre-election drought constructed from the monthly county-level Palmer Z Index; N_{it}^m is the bilateral-flow-weighted mean drought intensity among food-network partners, defined in equation (17); and N_{it}^{Δ} is the directional network drought exposure defined in equation (19). The dependent variable y_{it} is the Democratic two-party vote share in county i in presidential election year t.

The main identification concern is that drought exposure among a county's network partners, captured by N_{it}^m and N_{it}^{Δ}, may be correlated with unobserved political or economic shocks affecting the same broader region. In that case, the coefficients on the network variables and their interactions with historical Democratic vote share could capture regional co-movement rather than variation linked to pre-sample food-flow relationships.

The benchmark specification addresses this concern by including the 1,000 km regional drought measure, D_{it}^{reg}, defined in equation (22), which controls for surrounding drought conditions; county fixed effects, α_i, which absorb persistent differences across counties; election fixed effects, λ_t, which absorb national election-specific shocks; and the differential trend $\pi_i(t - \bar{t})$, which allows counties with different historical Democratic vote shares to follow different linear electoral trajectories over the 2008–2024 period. Conditional on these controls, the network coefficients are estimated from differences in partner-county drought exposure associated with counties' pre-sample food-flow connections, beyond their local and regional drought conditions and time-invariant characteristics.[7]

A separate concern arises from partisan realignment along the urban–rural divide. The sample period coincides with a widening gap in partisan voting between metropolitan and

[7]The regional and network measures differ in both their geographic support and their weighting scheme. D_{it}^{reg} assigns equal weight to all other counties within 1,000 km, whereas N_{it}^m assigns predetermined 2007 food-flow weights to economically connected counties, including partners that may lie inside or outside that radius. Conditional on local and regional drought, the network coefficient is therefore identified from differences between the drought distribution across a county's economically important partners and average drought in its broader geographic environment. The network estimate should not be interpreted as a comparison between purely local and purely distant shocks.

nonmetropolitan counties, with nonmetropolitan areas shifting increasingly toward the Republican Party (McKee, 2008; Scala and Johnson, 2017; Brown and Mettler, 2024). If this differential political evolution is correlated with both historical Democratic support and drought exposure, it could confound the estimated interactions between π_i and the drought measures.

We therefore consider, alongside the benchmark specification, an augmented specification that allows the partisan trend to differ between metropolitan and nonmetropolitan counties:

$$y_{it} = \alpha_i + \lambda_t + \sum_k X_{it}^k \left[\beta_{0k} + \beta_{1k}\pi_i\right] + \theta D_{it}^{reg} + \eta\,\pi_i\Big(t - \bar{t}\Big) + \zeta\,\pi_i R_i\Big(t - \bar{t}\Big) + \varepsilon_{it}. \qquad (24)$$

Here, R_i is an indicator equal to one for nonmetropolitan counties and zero for metropolitan counties. It is constructed from the 2003 Rural–Urban Continuum Codes (RUCC) as:

$$R_i = \begin{cases} 0, & \text{if } \mathrm{RUCC}_i^{2003} < 4 \quad \text{(metropolitan county)}, \\ 1, & \text{if } \mathrm{RUCC}_i^{2003} \geq 4 \quad \text{(nonmetropolitan county)}. \end{cases}$$

RUCC codes 1–3 identify metropolitan counties, whereas codes 4–9 identify nonmetropolitan counties according to urban population and adjacency to a metropolitan area. Using the 2003 classification fixes metropolitan status before the 2008–2024 estimation period. The additional term $\pi_i R_i(t - \bar{t})$ therefore allows the electoral trend associated with historical Democratic support to differ specifically in nonmetropolitan counties.

Equation (23) remains the benchmark specification used to characterize the main partisan gradients and sign-change thresholds. Equation (24) is used as an augmented specification to assess whether those relationships are sensitive to differential nonmetropolitan partisan realignment. Reporting the two specifications separately also makes transparent how much of the estimated heterogeneity is affected by this additional trend.

Beyond this differential-trend specification, we replace the common election fixed effects in the baseline specification successively with Census-region-by-election, state-by-election, and metropolitan-status-by-election fixed effects. These fixed effects absorb average election-specific changes common to counties within the same region, state, or metropolitan category. The coefficients on the drought variables are therefore estimated from variation across counties within the same classification and election, providing a stricter test of whether the baseline estimates reflect geographically concentrated election-specific shocks.

The coefficients of interest are β_{0k} and β_{1k}. A positive β_{1k} implies that the relationship between exposure X_{it}^k and Democratic vote share becomes more positive as historical Democratic vote share increases. When $\beta_{0k} < 0$, this relationship changes from negative to positive at $-\beta_{0k}/\beta_{1k}$, provided that the crossing point lies within the observed range of

π_i.

The empirical analysis evaluates the three predictions derived in Section 2.5: whether the relationship between drought and voting varies with historical Democratic vote share; whether local and network drought follow the same sign-changing pattern; and whether network drought operates through overall exposure among connected counties rather than through differences between suppliers and destination markets.

The network exposure measure has a shift-share form (Adão et al., 2019): it combines election-year drought in partner counties with bilateral food-flow weights fixed at their estimated 2007 values, before the first election in the sample. The weights therefore cannot adjust mechanically to subsequent drought conditions or electoral outcomes.

The empirical design therefore combines changes over time in the geographic distribution of drought with cross-county differences in pre-sample food-flow connections. The interaction with historical Democratic vote share, π_i, uses a measure of county political orientation constructed entirely from elections held before the estimation period. Thus, the specification examines whether drought exposure among food-flow partners is associated with electoral outcomes in connected counties and whether this relationship varies systematically with counties' historical Democratic vote share.

All reported standard errors are Conley-type spatial heteroskedasticity- and autocorrelation-consistent standard errors (Conley, 1999), calculated using a 500 km spatial cutoff and a Bartlett kernel, with serial correlation allowed over one presidential-election lag. This estimator allows the regression errors to be correlated across nearby counties and between consecutive presidential elections. Under the Bartlett kernel, the weight assigned to cross-county error covariance declines linearly with distance and reaches zero at the 500 km cutoff.

The choice of 500 km provides an intermediate benchmark relative to the geographic reach of the food-flow network. As shown in Table 2, the median flow-weighted partner distance is 265 km, while its 90th percentile is approximately 701 km. These network distances do not determine the appropriate range of residual spatial dependence, but they provide a useful reference for assessing the benchmark cutoff. We therefore examine the sensitivity of inference to a wider spatial range by re-estimating the Conley covariance matrix using an 800 km cutoff in Section 6.2.

5. Results

5.1 Baseline estimates

Table 3 reports the benchmark estimates. Each estimate is expressed as the percentage-point change in the Democratic two-party vote share associated with a one-unit increase

in the relevant drought measure.[8]

Table 3: Baseline estimates: local drought and network exposure

Drought-exposure channels	Local only D only	Local and network D, N^m	Local and asymmetric network D, N^Δ	All channels D, N^m, N^Δ
Local drought, D_{it}	-1.520^{***}	-0.843^{***}	-1.538^{***}	-0.879^{***}
	(0.310)	(0.321)	(0.305)	(0.317)
$D_{it} \times \pi_i$	3.662^{***}	2.135^{***}	3.681^{***}	2.174^{***}
	(0.740)	(0.799)	(0.731)	(0.791)
Mean network, N^m_{it}	–	-1.170^{***}	–	-1.135^{***}
	–	(0.407)	–	(0.396)
$N^m_{it} \times \pi_i$	–	2.608^{***}	–	2.596^{***}
	–	(0.910)	–	(0.885)
Asymmetric network, N^Δ_{it}	–	–	-0.421	-0.225
	–	–	(0.270)	(0.253)
$N^\Delta_{it} \times \pi_i$	–	–	0.364	-0.077
	–	–	(0.600)	(0.557)
Regional control, D^{reg}_{it}	0.593	0.633	0.586	0.604
	(0.476)	(0.506)	(0.475)	(0.505)
$\pi_i \times (t - \bar{t})$	-1.188^{***}	-1.162^{***}	-1.186^{***}	-1.162^{***}
	(0.190)	(0.192)	(0.190)	(0.192)
Observations	15 345	15 345	15 345	15 345
R^2	0.953	0.953	0.953	0.953
Joint Wald test of shock terms, W/q	12.771	7.241	8.350	6.177
p-value	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$

Note: The dependent variable is the Democratic two-party vote share. Conley standard errors using a 500 km spatial cutoff, a Bartlett kernel, and one presidential-election lag are reported in parentheses. Network exposure is constructed from 2007 SCTG 02–04 flows. The joint Wald test evaluates whether the exposure measures and their interactions with historical Democratic support included in each column are jointly equal to zero; W/q denotes the normalized Wald statistic. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Three patterns emerge. First, the estimated relationship between drought and the Democratic vote share varies with counties' historical Democratic vote share rather than shifting all counties uniformly toward one party. In column (1), the negative coefficient on local drought and the positive interaction with historical Democratic vote share imply that the estimated marginal relationship equals zero when historical Democratic vote share is $1.520/3.662 \simeq 0.415$, close to the sample mean of 0.448 (Table B.2). Local drought is therefore associated with a decline in Democratic vote share in counties below this threshold and an increase in counties above it. This pattern corresponds to the sign-changing partisan relationship derived in Section 2.5. At the sample mean, $\pi_i = 0.448$,

[8] Because local drought is standardized, one unit corresponds to one standard deviation of D_{it}. Mean network exposure is not standardized and has a sample standard deviation of 0.838.

the implied marginal relationship is approximately 0.12 percentage points in column (1), conditional on the other covariates and fixed effects.

Second, mean food-network drought exposure is statistically significant and also varies with historical Democratic vote share. The negative coefficient on N_{it}^{m} and its positive interaction with π_i reproduce the sign-changing pattern observed for local drought. The implied threshold is approximately $1.170/2.608 \simeq 0.449$ in column (2) and $1.135/2.596 \simeq 0.437$ in the full specification, both close to the sample mean of 0.448. Drought among food-flow partners is therefore associated with lower Democratic vote share in counties below the relevant threshold and higher Democratic vote share in counties above it, even after controlling for local drought and the broader regional drought environment.

The inclusion of N_{it}^{m} reduces the absolute magnitudes of both the local-drought coefficient and its interaction with historical Democratic vote share, from -1.520 and 3.662 in column (1) to -0.843 and 2.135 in column (2). This change does not imply that the marginal relationship with local drought declines by the same proportion at every value of π_i. At the sample mean, $\pi_i = 0.448$, the implied marginal relationship with local drought is approximately 0.11 percentage points in column (2) and 0.09 percentage points in column (4). Because the sample mean lies close to the estimated sign-change threshold, a one-standard-deviation increase in mean network exposure (0.838 units) has an implied relationship with Democratic vote share of approximately zero in column (2) and 0.02 percentage points in column (4). The network-exposure coefficients remain similar and statistically significant after controlling for regional drought, suggesting that broad regional correlation in drought does not fully account for the estimated network relationship.

Third, the estimates are more precise for overall drought exposure among food-flow partners than for directional differences between suppliers and destination markets. Neither directional network exposure, N_{it}^{Δ}, nor its interaction with historical Democratic vote share is statistically significant, and including these terms leaves the coefficients on local and mean network drought largely unchanged. The estimates therefore do not provide evidence of a distinct directional relationship operating through suppliers or destination markets. By contrast, the relationship with mean network exposure is precisely estimated.[9]

Finally, the joint Wald test rejects the null hypothesis that the included drought coefficients and interactions are jointly equal to zero in each of the four specifications. Taken together, the benchmark estimates show a sign-changing partisan relationship for both local drought and mean network exposure, while providing little evidence of an additional directional network relationship.

Because the sample period also coincides with substantial partisan realignment along

[9]Unlike mean network exposure, N_{it}^{Δ} is a signed difference whose weights do not sum to one. Its coefficient is therefore not directly comparable in magnitude with that on N_{it}^{m}. Moreover, supplier and destination-market exposures may offset within the directional index, reducing its variation and the power to detect a directional relationship.

the urban–rural divide, we next examine the augmented specification introduced in Equation (24). This specification adds the differential nonmetropolitan partisan trend $\pi_i R_i(t-\bar{t})$ while otherwise retaining the same controls, fixed effects, and exposure measures as the benchmark specification.

Table 4 reports the corresponding estimates. The coefficient on $\pi_i R_i(t-\bar{t})$ is approximately -0.94 in all four specifications and is statistically significant at the 1 percent level. Thus, among counties with higher historical Democratic support, Democratic vote share declined more rapidly over 2008–2024 in nonmetropolitan counties than in metropolitan counties, consistent with the partisan realignment described in Section 4.

Table 4: Augmented estimates with differential nonmetropolitan partisan trends

Drought-exposure channels	Local only D only	Local and network D, N^m	Local and asymmetric network D, N^Δ	All channels D, N^m, N^Δ
Local drought, D_{it}	-1.341^{***}	-0.634^{**}	-1.350^{***}	-0.682^{**}
	(0.282)	(0.285)	(0.278)	(0.281)
$D_{it} \times \pi_i$	3.251^{***}	1.634^{**}	3.250^{***}	1.704^{**}
	(0.675)	(0.713)	(0.667)	(0.708)
Mean network exposure, N^m_{it}	–	-1.209^{***}	–	-1.140^{***}
	–	(0.390)	–	(0.379)
$N^m_{it} \times \pi_i$	–	2.761^{***}	–	2.663^{***}
	–	(0.884)	–	(0.861)
Asymmetric network exposure, N^Δ_{it}	–	–	-0.578^{**}	-0.379
	–	–	(0.260)	(0.242)
$N^\Delta_{it} \times \pi_i$	–	–	0.822	0.370
	–	–	(0.589)	(0.548)
Regional control, D^{reg}_{it}	0.515	0.539	0.513	0.517
	(0.449)	(0.476)	(0.448)	(0.476)
$\pi_i \times (t-\bar{t})$	-0.660^{***}	-0.632^{***}	-0.658^{***}	-0.633^{***}
	(0.179)	(0.181)	(0.178)	(0.181)
$\pi_i \times R_i \times (t-\bar{t})$	-0.940^{***}	-0.941^{***}	-0.939^{***}	-0.940^{***}
	(0.057)	(0.057)	(0.057)	(0.057)
Observations	15 345	15 345	15 345	15 345
R^2	0.959	0.959	0.959	0.959
Joint Wald test of shock terms, W/q	12.106	7.031	7.604	5.547
p-value	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$

Note: The dependent variable is the Democratic two-party vote share. All specifications include county and election fixed effects and the differential nonmetropolitan partisan trend $\pi_i R_i(t-\bar{t})$, where $R_i = 1$ for nonmetropolitan counties under the 2003 RUCC classification. Conley standard errors using a 500 km spatial cutoff, a Bartlett kernel, and one presidential-election lag are reported in parentheses. The joint Wald test evaluates whether the included drought-exposure measures and their interactions with historical Democratic support are jointly equal to zero; W/q denotes the normalized Wald statistic. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Allowing for this differential trend reduces the estimated local-drought gradient. Relative to the corresponding benchmark estimates, the interaction between local drought and historical Democratic support declines from 3.662 to 3.251 in the drought-only specification and from 2.174 to 1.704 in the all-channels specification. The interaction nevertheless remains positive and statistically significant. By contrast, the mean-network interaction is highly stable: in the all-channels specification it changes from 2.596 to 2.663 and remains statistically significant at the 1 percent level. Neither the directional-network interaction nor the directional level coefficient is statistically significant in the all-channels specification, and the drought-exposure terms remain jointly significant in every specification ($p < 0.001$).

These comparisons indicate that differential partisan trends in nonmetropolitan counties account for part of the estimated heterogeneity in the local-drought relationship but do not explain the mean-network relationship. They also reinforce the distinction between the two results: the network gradient is particularly stable to accounting for rural partisan realignment, whereas the local gradient is more sensitive in magnitude.

5.2 Heterogeneity by historical Democratic vote share

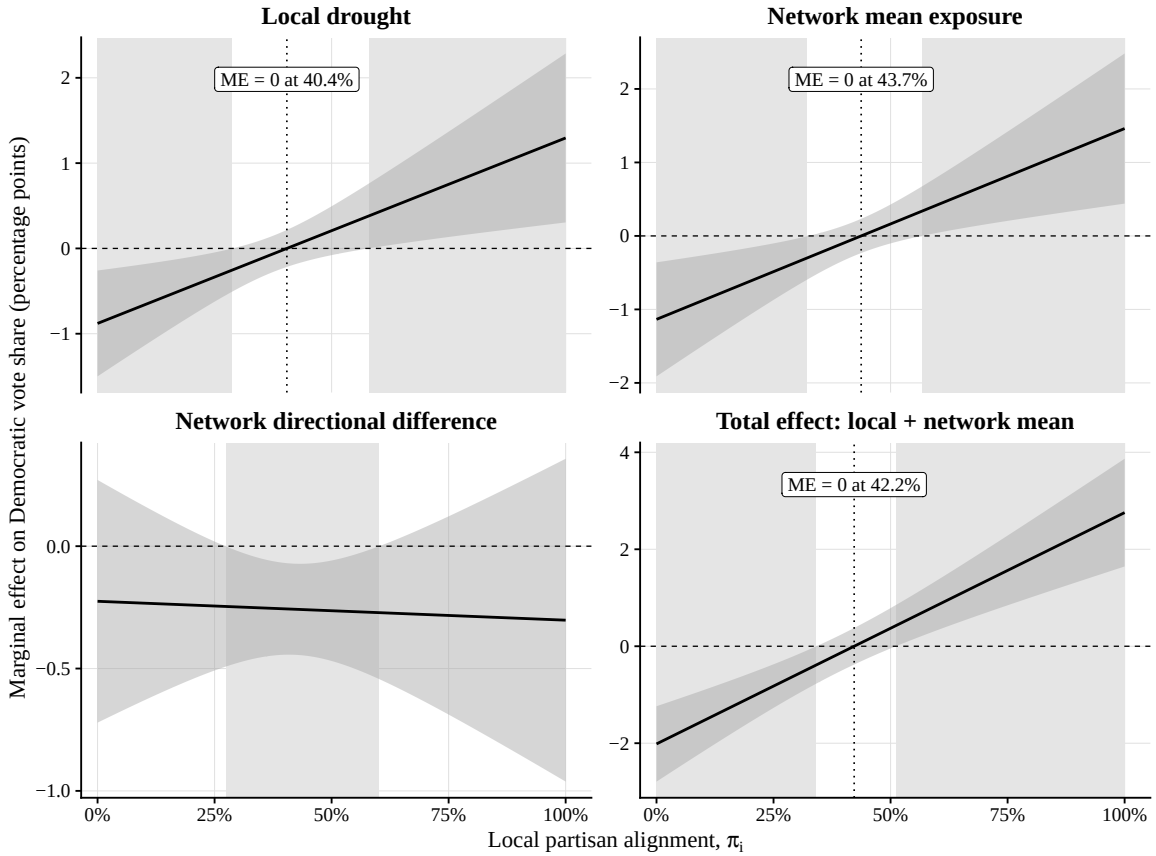
Because the empirical specification interacts drought exposure with historical Democratic vote share, the coefficients on the exposure variables correspond to their estimated marginal effects when $\pi_i = 0$. Because this value lies outside the observed range of π_i, these coefficients represent extrapolations and are not directly informative about a typical county. We therefore evaluate the marginal effects at values representative of historically Republican counties, counties with balanced historical vote shares, and historically Democratic counties.

To characterize the sign-changing relationships and evaluate the restrictions derived in Section 2.5, Figure 2 and Table 5 report results from the all-channels benchmark specification. The figure traces the marginal effects across the distribution of π_i, while the table reports estimates and Conley standard errors at selected values of π_i.

The results show a clear sign reversal in the estimated marginal relationship between local drought and Democratic vote share across levels of historical Democratic support. In the all-channels benchmark specification, the estimated effect is negative at low values of π_i, close to zero near the center of the distribution, and positive at high values. At $\pi_i = 0.25$, a one-standard-deviation increase in local drought is associated with a 0.336-percentage-point decrease in the Democratic vote share, and the estimate is statistically significant at the 5 percent level. At $\pi_i = 0.50$, the estimated effect is small and statistically indistinguishable from zero. At $\pi_i = 0.75$, a one-standard-deviation increase in local drought is associated with a 0.751-percentage-point increase in the Democratic vote share, with the estimate again statistically significant at the 5 percent level. The estimated marginal effect equals

zero at $\pi_i \simeq 0.404$, as shown in the upper-left panel of Figure 2. Thus, the relationship between local drought and the Democratic vote share varies systematically with historical Democratic support and changes sign within its observed range.

Figure 2: Marginal effects of drought by historical Democratic vote share



Note: The figure reports marginal effects from Equation (23). Confidence bands are based on Conley standard errors using a 500 km spatial cutoff, a Bartlett kernel, and one presidential-election lag.

Table 5: Marginal effects of drought at selected levels of historical Democratic vote share

Drought-exposure channels	Marginal effect	Historical Democratic support, π_i				
		$0.00^{\dagger}$	0.25	0.50	0.75	$1.00^{\dagger}$
D only	Local drought, D_{it}	-1.520^{***}	-0.605^{***}	0.311^{*}	1.226^{***}	2.142^{***}
		(0.310)	(0.164)	(0.161)	(0.306)	(0.479)
D, N^m	Local drought, D_{it}	-0.843^{***}	-0.309^{**}	0.224	0.758^{**}	1.292^{**}
		(0.321)	(0.150)	(0.147)	(0.317)	(0.509)
	Mean network exposure, N^m_{it}	-1.170^{***}	-0.518^{***}	0.134	0.786^{**}	1.438^{***}
		(0.407)	(0.200)	(0.132)	(0.314)	(0.532)
	Total, D_{it}, N^m_{it}	-2.013^{***}	-0.828^{***}	0.358^{*}	1.544^{***}	2.729^{***}
		(0.404)	(0.232)	(0.211)	(0.368)	(0.572)
D, N^{Δ}	Local drought, D_{it}	-1.538^{***}	-0.618^{***}	0.302^{*}	1.223^{***}	2.143^{***}
		(0.305)	(0.161)	(0.160)	(0.303)	(0.475)
	Asymmetric network exposure, N^{Δ}_{it}	-0.421	-0.330^{**}	-0.239^{**}	-0.148	-0.057
		(0.270)	(0.140)	(0.105)	(0.218)	(0.359)
D, N^m, N^{Δ}	Local drought, D_{it}	-0.879^{***}	-0.336^{**}	0.208	0.751^{**}	1.295^{**}
		(0.317)	(0.147)	(0.146)	(0.315)	(0.505)
	Mean network exposure, N^m_{it}	-1.135^{***}	-0.486^{**}	0.163	0.812^{***}	1.461^{***}
		(0.396)	(0.196)	(0.135)	(0.310)	(0.521)
	Asymmetric network exposure, N^{Δ}_{it}	-0.225	-0.244^{*}	-0.264^{**}	-0.283	-0.302
		(0.253)	(0.134)	(0.105)	(0.207)	(0.337)
	Total, D_{it}, N^m_{it}	-2.015^{***}	-0.822^{***}	0.371^{*}	1.563^{***}	2.756^{***}
		(0.397)	(0.229)	(0.212)	(0.367)	(0.567)

Note: The table reports marginal effects evaluated at the indicated values of π_i, with Conley standard errors in parentheses. The "Total" rows correspond to simultaneous one-unit increases in local drought and mean network exposure. Values marked with † lie outside the observed range, $\pi_i \in [0.12, 0.835]$, and are linear extrapolations. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Mean drought exposure through the food network exhibits a similar pattern. In the all-channels benchmark specification, the estimated marginal effect of a one-unit increase in N_{it}^{m} is -0.486 percentage points at $\pi_i = 0.25$, is close to zero at $\pi_i = 0.50$, and rises to 0.812 percentage points at $\pi_i = 0.75$. The estimates at $\pi_i = 0.25$ and $\pi_i = 0.75$ are statistically significant at the 5 and 1 percent levels, respectively. Because the sample standard deviation of mean network exposure is 0.838, these estimates correspond to approximately -0.41 and $+0.68$ percentage points for a one-standard-deviation increase in network exposure. The estimated marginal effect equals zero at $\pi_i \simeq 0.437$, only 0.033 above the corresponding threshold for local drought. A Wald test of the restriction that the local- and mean-network relationships have the same zero-crossing point does not reject equality. The estimated thresholds are therefore consistent with the common-scaling restriction derived in Section 2.5.

The estimated joint effect of simultaneous one-unit increases in local and mean network drought is larger in absolute magnitude, especially at low and high values of historical Democratic support, and exhibits the same sign reversal. At $\pi_i = 0.25$, simultaneous increases in both exposure measures are associated with a 0.822-percentage-point decrease in the Democratic vote share. At $\pi_i = 0.50$, the estimated joint effect is 0.371 percentage points and is statistically significant at the 10 percent level. At $\pi_i = 0.75$, it reaches 1.563 percentage points and is statistically significant at the 1 percent level. The estimated joint effect equals zero at $\pi_i \simeq 0.422$, between the zero-crossing points for local and network exposure. The two component relationships reinforce one another over most of the distribution but have opposite signs between the two thresholds. The joint relationship is small near the center of the distribution because both components are close to zero and partly offset one another in this interval.

By contrast, directional network exposure does not vary systematically with historical Democratic support. Its estimated marginal effect is negative across the selected values of π_i, but the confidence interval in Figure 2 includes zero over most of the distribution. Although Table 5 shows statistical significance at some intermediate values of π_i, neither the coefficient on directional network exposure nor its interaction with historical Democratic support is statistically significant in the all-channels benchmark specification. The results therefore provide little evidence that the relationship between network drought and electoral outcomes depends on whether exposed counties are primarily suppliers or destination markets. The evidence is stronger for mean network exposure, which displays a clear sign-changing pattern across historical Democratic vote shares.

Table 5 also shows that the qualitative pattern is stable across the alternative benchmark model specifications. Estimating the local-drought channel alone produces the same upward pattern in the marginal effect as historical Democratic support increases. Adding mean network exposure reduces the magnitude of this gradient but does not remove it. Likewise, the estimated marginal effect of mean network exposure remains negative at

low values of π_i and positive at high values after the directional network component is added. The sign reversals for local drought and mean network exposure are therefore not specific to the all-channels specification.

Taken together, Figure 2 and Table 5 show that the estimated relationships for both local and mean network drought change sign within the observed distribution of historical Democratic support. The proximity of the two zero crossings is also preserved after accounting for differential nonmetropolitan partisan trends.

5.3 Network exposure when the local drought measure is zero

The baseline specification estimates the relationship between the Democratic two-party vote share and mean network drought exposure, N_{it}^m. This measure applies to two distinct types of county-election observations: counties that experience local drought while also being exposed through their food-flow partners, and counties for which the local drought measure equals zero but whose food-flow partners experience drought. The second group helps isolate the network pattern because it allows us to examine whether drought occurring outside the county is associated with the Democratic two-party vote share when local drought intensity is measured as zero. To make this distinction explicit, we construct a set of mutually exclusive exposure categories.

The presence of local drought is defined using $\widetilde{D}_{it} = \max\{-Z_{it}, 0\}$, introduced in Section 3.2, where Z_{it} is the day-weighted average Palmer Z Index for county i over the six months ending in September of election year t. The standardized measure satisfies $D_{it} = 0$ when $Z_{it} \geq 0$ and is positive when the six-month average indicates a moisture deficit. We define $L_{it} = \mathbf{1}\{D_{it} > 0\}$, where $L_{it} = 1$ indicates that the county's average moisture conditions over the six-month pre-election period show a deficit.

Network exposure is classified using the sample median of mean network drought exposure, $q_{0.50}(N^m) = 0.347$. A county-election observation is assigned to the high-network-exposure category when N_{it}^m exceeds this value. We therefore define $H_{it} = \mathbf{1}\{N_{it}^m > q_{0.50}(N^m)\}$ where $q_{0.50}(N^m)$ denotes the median of N_{it}^m in the estimation sample. Observations at or below the median are assigned to the lower-exposure category.

The main typology combines the indicator for positive local drought intensity, L_{it}, with the indicator for above-median network exposure, H_{it}, to form four mutually exclusive categories: (i) unaffected: $L_{it} = 0$, $H_{it} = 0$; (ii) local only: $L_{it} = 1$, $H_{it} = 0$; (iii) network only: $L_{it} = 0$, $H_{it} = 1$; and (iv) local and network: $L_{it} = 1$, $H_{it} = 1$.

The unaffected category serves as the reference group. The main coefficient of interest is the interaction between the network-only indicator and historical Democratic vote share. This coefficient shows whether the difference in Democratic vote share between the network-only group and the reference group becomes larger or smaller as historical Democratic vote share increases. The network-only group consists of county-election

observations with no measured local drought, $D_{it} = 0$, but above-median drought exposure among food-flow partners, $H_{it} = 1$.

This comparison has two limitations. First, the reference group may include counties with positive network exposure close to the median, so the contrast with the above-median group may be relatively small. Second, the median split reduces a continuous measure to two categories and may therefore hide how the relationship changes across different levels of network exposure.

We therefore construct a more detailed classification based on network-exposure intensity. Among observations with no measured local drought, network exposure is divided into three mutually exclusive categories: (i) very low network exposure (reference), $D_{it} = 0$, $N_{it}^m \leq q_{0.25}$; (ii) low network exposure, $D_{it} = 0$, $q_{0.25} < N_{it}^m \leq q_{0.50}$; and (iii) high network exposure, $D_{it} = 0$, $N_{it}^m > q_{0.50}$. Here, $q_{0.25}$ and $q_{0.50}$ denote the 25th and 50th percentiles of N_{it}^m in the full estimation sample. These categories are complemented by two groups with positive local drought: (iv) local drought with network exposure at or below the median, $D_{it} > 0, N_{it}^m \leq q_{0.50}$; and (v) local drought with above-median network exposure, $D_{it} > 0, N_{it}^m > q_{0.50}$.

This refinement allows us to assess whether the relationship between network drought exposure and the Democratic two-party vote share becomes more pronounced as exposure increases. Under the network-channel interpretation, differences across counties with lower and higher historical Democratic vote shares should be relatively small in the low-exposure category and larger in the high-exposure category. In particular, the difference in Democratic vote share between the low- and very-low-exposure categories should be smaller than the difference between the high- and very-low-exposure categories.

For both classifications, let G_{it}^g denote an indicator equal to one when county-election observation (i, t) belongs to exposure category g, and zero otherwise. The set of categories depends on the classification being considered. In the main typology, g indexes the three non-reference categories "local only," "network only," and "local and network," with the "unaffected" category omitted. In the more detailed typology, g indexes "low network," "high network," "local drought with network exposure at or below the median," and "local drought with above-median network exposure," with the "very low network" category among observations with no measured local drought omitted.

We then estimate:

$$y_{it} = \alpha_i + \lambda_t + \sum_{g \neq g_0} G_{it}^g \Big[\gamma_{0g} + \gamma_{1g} \pi_i \Big] + \theta D_{it}^{\text{reg}} + \eta \pi_i \Big(t - \bar{t} \Big) + \zeta \pi_i R_i \Big(t - \bar{t} \Big) + \varepsilon_{it}. \quad (25)$$

where g_0 denotes the omitted reference category.

For each non-reference exposure category g, γ_{0g} measures the difference in Democratic vote share between category g and the omitted reference category when $\pi_i = 0$, while γ_{1g} measures how this category-reference difference varies with historical Democratic support.

Accordingly, at a given value of π_i, the estimated difference between category g and the reference category is $\gamma_{0g} + \gamma_{1g}\pi_i$.

For each typology, we report two specifications. The benchmark specification sets $\zeta = 0$, whereas the augmented specification includes $\pi_i R_i(t - \bar{t})$, allowing the relationship between historical Democratic support and the Democratic two-party vote share to follow a different trend in nonmetropolitan counties.

Table 6 describes the composition of the exposure categories.

Table 6: Composition of local and network drought-exposure categories

Category	County-elections	Share	Counties	States	$\bar{\pi}_i$	$\bar{N}^m$
Panel A — Main typology ($L_{it} \times H_{it}$)						
Unaffected	5733	37.4 %	2520	47	0.458	0.076
Local only	1940	12.6 %	1369	47	0.452	0.133
Network only	1663	10.8 %	1230	47	0.448	0.820
Local and network	6009	39.2 %	2818	47	0.438	1.441
Panel B — Network-exposure intensity gradient (counties without local drought)						
No local, very low network	3189	20.8 %	1820	42	0.461	0.013
No local, low network	2544	16.6 %	1657	46	0.453	0.156
No local, high network	1663	10.8 %	1230	47	0.448	0.820
Local, low or no network	1940	12.6 %	1369	47	0.452	0.133
Local, high network	6009	39.2 %	2818	47	0.438	1.441

Notes: Panel A crosses positive local drought with above-median mean network exposure, where $q_{0.50}(N^m) = 0.347$. Panel B further divides observations with no local drought using the 25th and 50th percentiles of mean network exposure, computed over the full estimation sample ($N = 15{,}345$).

The first panel reports the four-category typology, while the second presents the more detailed classification by network-exposure intensity. The categories contain substantial numbers of county-election observations, include counties from many states, and have relatively similar average levels of historical Democratic vote share. These patterns reduce concerns that the results are driven by a small number of observations concentrated in a few locations or at unusually high or low values of historical Democratic vote share.

Table 7 reports the category-based regressions for the two exposure classifications, both with and without the differential nonmetropolitan partisan trend.

Table 7: Differences in Democratic vote share by local and network drought-exposure category

Specification	Benchmark (1)	+ Rural trend (2)
Panel A — Main typology (reference: unaffected)		
Local only	-2.277^{***}	-1.718^{**}
	(0.794)	(0.742)
Local only $\times\pi_i$	4.758^{***}	3.630^{**}
	(1.815)	(1.725)
Network only	-1.970^{***}	-1.764^{***}
	(0.589)	(0.595)
Network only $\times\pi_i$	3.869^{***}	3.538^{***}
	(1.315)	(1.324)
Local and network	-4.277^{***}	-3.456^{***}
	(0.793)	(0.766)
Local and network $\times\pi_i$	9.419^{***}	7.697^{***}
	(1.827)	(1.770)
Panel B — Cross-classification by local drought and network-exposure level (reference: no local drought, very low network exposure)		
No local, low network	-1.006^{*}	-0.818
	(0.590)	(0.573)
No local, low network $\times\pi_i$	1.467	0.988
	(1.218)	(1.198)
No local, high network	-2.452^{***}	-2.160^{***}
	(0.673)	(0.694)
No local, high network $\times\pi_i$	4.559^{***}	4.007^{***}
	(1.438)	(1.506)
Local, low or no network	-2.707^{***}	-2.079^{**}
	(0.856)	(0.821)
Local, low or no network $\times\pi_i$	5.380^{***}	4.076^{**}
	(1.901)	(1.855)
Local, high network	-4.765^{***}	-3.864^{***}
	(0.860)	(0.849)
Local, high network $\times\pi_i$	10.120^{***}	8.192^{***}
	(1.906)	(1.914)
Common controls		
Regional control, D_{it}^{reg}	0.736^{*} / 0.749^{*}	0.630 / 0.642
$\pi_i \times (t-\bar{t})$	-1.180^{***} / -1.178^{***}	-0.659^{***} / -0.658^{***}
$\pi_i \times R_i \times (t-\bar{t})$	—	-0.935^{***}
Observations	15,345	15,345
R^2	0.953	0.959
Normalized joint Wald statistic for cell terms, W/q	5.082 / 4.098	3.617 / 3.173
p-values	$p < 0.001$ / $p < 0.001$	$p = 0.001$ / $p = 0.001$

Notes: The dependent variable is the Democratic two-party vote share. All specifications include county and election fixed effects; column (2) additionally includes the differential nonmetropolitan partisan trend, $\pi_i R_i(t-\bar{t})$. Conley standard errors using a 500 km spatial cutoff, a Bartlett kernel, and one presidential-election lag are reported in parentheses. The joint Wald tests evaluate the non-reference category coefficients and their interactions with historical Democratic support ($q = 6$ in Panel A and $q = 8$ in Panel B); W/q denotes the normalized Wald statistic. When two values are reported as A/B, they correspond to Panels A and B, respectively. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Panel A presents the main four-category typology. Relative to county-election observations with no local drought and below-median network exposure, the "network only" category has a level coefficient of -1.970 and an interaction coefficient of 3.869 in the benchmark specification. Both are statistically significant at the 1 percent level. These estimates are somewhat smaller than the corresponding coefficients for the "local only" category, -2.277 and 4.758, but they have the same signs. This result shows that above-median drought exposure among food-flow partners is related to Democratic vote share even when the local drought measure equals zero, and that this relationship varies with historical Democratic vote share.

The pattern remains after adding the differential nonmetropolitan partisan trend. The interaction for the "network only" category declines from 3.869 to 3.538 but remains statistically significant at the 1 percent level. The interaction therefore remains positive and precisely estimated after allowing historical Democratic vote share to follow a different trend in nonmetropolitan counties.

The coefficients for the "local and network" category are broadly consistent with an additive pattern. In the benchmark specification, its level coefficient, -4.277, is close to the sum of the "local only" and "network only" coefficients: $-2.277 - 1.970 = -4.247$. Its interaction coefficient, 9.419, is somewhat larger than the corresponding sum: $4.758 + 3.869 = 8.627$. The same broad pattern holds after adding the nonmetropolitan trend. This comparison is descriptive rather than a formal test of additivity, because the categories differ in their average drought intensity and the coefficient restrictions are not directly tested.

Panel B of Table 7 examines whether the network relationship differs by exposure intensity among observations with no local drought. The interaction coefficient is 1.467 for the low-network category and 4.559 for the high-network category in the benchmark specification. After adding the nonmetropolitan trend, the corresponding estimates are 0.988 and 4.007. The low-exposure interactions are not statistically significant, whereas the high-exposure interactions are statistically significant at the 1 percent level.

Table 8 translates the coefficients from Table 7 into estimated differences in Democratic vote share at selected values of π_i.

In the main typology, the estimated difference between the "network only" category and the reference category is -1.003 percentage points at $\pi_i = 0.25$, approximately zero at $\pi_i = 0.50$, and 0.932 percentage points at $\pi_i = 0.75$ in the benchmark specification. With the differential nonmetropolitan trend, the corresponding estimates are -0.880, 0.004, and 0.889 percentage points. Thus, the network-only difference is negative at low historical Democratic vote shares, close to zero near the center of the distribution, and positive at higher values.

Table 8: Estimated differences in Democratic vote share across local and network drought-exposure categories, by historical Democratic vote share

	Historical Democratic support, π_i				
Category	$0.00^{\dagger}$	0.25	0.50	0.75	$1.00^{\dagger}$
Panel A — Main exposure typology: benchmark specification					
Local only	-2.277^{***}	-1.088^{***}	0.102	1.292^{*}	2.481^{**}
	(0.794)	(0.394)	(0.305)	(0.665)	(1.097)
Network only	-1.970^{***}	-1.003^{***}	-0.036	0.932^{*}	1.899^{**}
	(0.589)	(0.314)	(0.258)	(0.501)	(0.807)
Local and network	-4.277^{***}	-1.922^{***}	0.433	2.788^{***}	5.143^{***}
	(0.793)	(0.425)	(0.386)	(0.731)	(1.157)
Panel B — Main exposure typology: benchmark specification + rural trend					
Local only	-1.718^{**}	-0.811^{**}	0.097	1.004	1.912^{*}
	(0.742)	(0.367)	(0.302)	(0.648)	(1.058)
Network only	-1.764^{***}	-0.880^{**}	0.004	0.889^{*}	1.773^{**}
	(0.595)	(0.313)	(0.247)	(0.493)	(0.802)
Local and network	-3.456^{***}	-1.531^{***}	0.393	2.317^{***}	4.242^{***}
	(0.766)	(0.408)	(0.371)	(0.707)	(1.120)
Panel C — Cross-classification by local drought and network-exposure level: benchmark (reference: no local drought and very low network exposure)					
No local, low network	-1.006^{*}	-0.639^{*}	-0.273	0.094	0.461
	(0.590)	(0.334)	(0.247)	(0.443)	(0.719)
No local, high network	-2.452^{***}	-1.312^{***}	-0.172	0.968^{*}	2.108^{**}
	(0.673)	(0.385)	(0.319)	(0.561)	(0.886)
Local, low or no network	-2.707^{***}	-1.361^{***}	-0.016	1.329^{*}	2.674^{**}
	(0.856)	(0.457)	(0.370)	(0.719)	(1.161)
Local, high network	-4.765^{***}	-2.235^{***}	0.295	2.825^{***}	5.355^{***}
	(0.860)	(0.488)	(0.437)	(0.773)	(1.208)
Panel D — Cross-classification by local drought and network-exposure level: rural trend (reference: no local drought and very low network exposure)					
No local, low network	-0.818	-0.571^{*}	-0.324	-0.077	0.170
	(0.573)	(0.319)	(0.233)	(0.432)	(0.706)
No local, high network	-2.160^{***}	-1.158^{***}	-0.157	0.845	1.847^{**}
	(0.694)	(0.384)	(0.310)	(0.573)	(0.919)
Local, low or no network	-2.079^{**}	-1.060^{**}	-0.041	0.978	1.997^{*}
	(0.821)	(0.435)	(0.367)	(0.714)	(1.147)
Local, high network	-3.864^{***}	-1.816^{***}	0.232	2.280^{***}	4.328^{***}
	(0.849)	(0.473)	(0.429)	(0.776)	(1.216)

Notes: Each cell reports $\gamma_{0g} + \gamma_{1g}\pi_i$ from the regressions in Table 7. Panels A and B use no local drought and below-median network exposure as the reference category; Panels C and D use no local drought and very low network exposure. Panels B and D include the differential nonmetropolitan partisan trend. Conley standard errors are reported in parentheses. Values marked with † lie outside the observed range, $\pi_i \in [0.12, 0.835]$, and are linear extrapolations. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The "local only" and "local and network" categories follow a similar sign-changing pattern. In the benchmark specification, their estimated differences change sign at approximately $\pi_i = 0.48$ and $\pi_i = 0.45$, respectively, while the network-only difference changes sign at approximately $\pi_i = 0.51$. These results show that the category-based estimates reproduce the sign reversal found in the continuous specification.

The more detailed classification in Panels C and D of Table 8 clarifies the role of network-exposure intensity. Among observations with no local drought, the low-network category differs little from the very-low-network reference category over much of the observed range of π_i. By contrast, the high-network category is associated with a difference of -1.312 percentage points at $\pi_i = 0.25$, -0.172 at $\pi_i = 0.50$, and 0.968 at $\pi_i = 0.75$ in the benchmark specification. The corresponding estimates with the nonmetropolitan trend are -1.158, -0.157, and 0.845 percentage points. The estimated differences are therefore larger for high network exposure than for low network exposure, especially at low and high values of historical Democratic vote share.

Taken together, the two tables provide consistent evidence that network drought is related to Democratic vote share even when the local drought measure equals zero and that this relationship is more pronounced at higher levels of network exposure.

6. Robustness and Sensitivity

6.1 Controlling for alternative election-specific fixed effects

The regional control, D_{it}^{reg}, accounts for surrounding drought conditions, while the election fixed effects, λ_t, absorb national political shocks common to a given election year. A remaining concern is that election-specific shocks affecting a particular region, state, or metropolitan category may be correlated with both drought exposure and historical Democratic vote share.

To address this possibility, we replace the common election fixed effects with Census-region-by-election, state-by-election, or metropolitan-status-by-election fixed effects. The fixed-effects component then becomes $\alpha_i + \lambda_{r(i)t}$, where $r(i)$ denotes county i's Census region, state, or metropolitan status, depending on the specification. These fixed effects absorb average election-specific shocks shared by counties within the same classification. The coefficients on local and network drought are therefore estimated using differences across counties within the same Census region, state, or metropolitan category and election year.

For each set of election-specific fixed effects, Table 9 reports both the benchmark specification and an augmented specification that also includes the differential trend for nonmetropolitan counties, $\pi_i R_i(t - \bar{t})$. The table therefore shows whether the estimated relationships between drought exposure and electoral outcomes remain robust when

election-specific shocks are absorbed at the Census-region, state, or metropolitan-status level.

Table 9: Robustness to region-, state-, and metropolitan-status-by-election fixed effects

	Region × election		State × election		Metro/nonmetro × election	
Specification	Benchmark	+ rural trend	Benchmark	+ rural trend	Benchmark	+ rural trend
Local drought, D_{it}	-0.799^{***}	-0.635^{**}	-0.596^{**}	-0.442^{*}	-0.656^{**}	-0.657^{**}
	(0.297)	(0.268)	(0.254)	(0.226)	(0.282)	(0.281)
$D_{it} \times \pi_i$	1.889^{**}	1.506^{**}	1.419^{**}	1.085^{**}	1.646^{**}	1.639^{**}
	(0.742)	(0.671)	(0.579)	(0.510)	(0.714)	(0.709)
Mean network, N_{it}^{m}	-1.086^{***}	-1.101^{***}	-0.999^{***}	-0.978^{***}	-1.168^{***}	-1.164^{***}
	(0.390)	(0.374)	(0.328)	(0.305)	(0.385)	(0.383)
$N_{it}^{m} \times \pi_i$	2.339^{***}	2.456^{***}	2.275^{***}	2.341^{***}	2.788^{***}	2.782^{***}
	(0.861)	(0.833)	(0.741)	(0.697)	(0.873)	(0.869)
Asymmetric network, N_{it}^{Δ}	-0.118	-0.253	-0.054	-0.184	-0.337	-0.345
	(0.245)	(0.233)	(0.208)	(0.204)	(0.234)	(0.235)
$N_{it}^{\Delta} \times \pi_i$	-0.167	0.220	-0.286	0.085	0.323	0.348
	(0.557)	(0.547)	(0.465)	(0.474)	(0.533)	(0.534)
Regional control, D_{it}^{reg}	-0.127	-0.153	-0.146	-0.142	0.367	0.364
	(0.349)	(0.334)	(0.503)	(0.485)	(0.478)	(0.478)
$\pi_i \times (t - \bar{t})$	-1.340^{***}	-0.822^{***}	-1.509^{***}	-0.964^{***}	-1.284^{***}	-1.085^{***}
	(0.184)	(0.174)	(0.174)	(0.156)	(0.186)	(0.210)
$\pi_i \times R_i \times (t - \bar{t})$	–	-0.869^{***}	–	-0.837^{***}	–	-0.289
	–	(0.049)	–	(0.046)	–	(0.212)
Observations	15 345	15 345	15 345	15 345	15 345	15 345
R^2	0.962	0.966	0.971	0.975	0.959	0.960
Joint Wald test of shock terms, W/q	5.370	4.919	4.646	4.232	5.182	5.192
p-value	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$	$p < 0.001$

Note: The dependent variable is the Democratic two-party vote share. Column pairs include Census-region-, state-, or metropolitan-status-by- election fixed effects. The + rural-trend specifications additionally include $\pi_i R_i (t - \bar{t})$. Conley standard errors are reported in parentheses. The joint Wald test evaluates whether the included drought-exposure measures and their interactions with historical Democratic support are jointly equal to zero; W/q denotes the normalized Wald statistic. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Relative to the baseline estimate reported earlier, the interaction between local drought and historical Democratic vote share becomes smaller when Census-region-by-election and state-by-election fixed effects absorb election-specific shocks common to counties within the same region or state. The coefficient is 1.89 under region-by-election fixed effects and 1.42 under state-by-election fixed effects, and remains statistically significant in both specifications. Under metropolitan-status-by-election fixed effects, the coefficient is 1.65 and is also statistically significant. Adding the differential nonmetropolitan partisan trend further reduces the interaction under the region and state specifications, to 1.51 and 1.09, respectively, but has almost no effect under the metropolitan-status specification.

The estimates associated with mean network drought exposure are more stable across fixed-effects structures. The interaction between mean network exposure and historical Democratic vote share ranges from 2.28 to 2.79 in the benchmark specifications and remains statistically significant at the 1 percent level under every fixed-effects structure. The estimates also change little when the differential nonmetropolitan trend is included. These results suggest that the relationship between mean network drought and Democratic vote share continues to vary with historical Democratic vote share after absorbing election-specific shocks common to counties within regions, states, or metropolitan categories.

By contrast, neither directional network exposure nor its interaction with historical Democratic vote share is statistically significant in any specification. The table therefore provides little evidence of an additional relationship based on the supplier–buyer composition of network exposure once mean network drought is included. Finally, the drought-related coefficients are jointly statistically significant in every specification ($p < 0.001$).

6.2 Network structure and spatial inference

The construction of network drought exposure gives greater weight to counties connected by larger bilateral food flows. As shown in Table 2, the typical county is connected to many external partners, but bilateral flows are relatively concentrated among a smaller number of economically important connections. This raises two related concerns. First, the estimated network relationship could be driven disproportionately by counties with exceptionally large total food flows, which occupy particularly important positions in the food-flow network. Second, counties whose bilateral weights are highly concentrated among a small number of partners may generate exposure measures that are especially sensitive to drought in those partners. We therefore examine whether the benchmark results are robust to excluding network hubs and counties with highly concentrated partner weights.

Specifically, column (2) of Table 10 excludes the 31 counties in the top 1 percent of the 2007 distribution of total external bilateral SCTG 02–04 food flows. Column (4)

instead excludes the 307 counties in the top decile of the Herfindahl index of bilateral partner weights. Network weights are not recomputed after these exclusions, and the exposure variables retain the normalization used in the full benchmark sample. This ensures that changes in the estimated coefficients reflect the sample restriction rather than a redefinition of network exposure.

Table 10: Robustness to network structure and spatial inference

	Benchmark 500 km (1)	Excluding top 1% flow (2)	Conley 800 km (3)	Excluding top HHI decile (4)
Local drought, D_{it}	-0.879^{***}	-0.872^{***}	-0.879^{***}	-0.871^{***}
	(0.317)	(0.318)	(0.309)	(0.318)
$D_{it} \times \pi_i$	2.174^{***}	2.158^{***}	2.174^{***}	2.148^{***}
	(0.791)	(0.794)	(0.766)	(0.800)
Mean network exposure, N_{it}^{m}	-1.135^{***}	-1.142^{***}	-1.135^{***}	-1.328^{***}
	(0.396)	(0.396)	(0.401)	(0.444)
$N_{it}^{m} \times \pi_i$	2.596^{***}	2.637^{***}	2.596^{***}	3.126^{***}
	(0.885)	(0.886)	(0.902)	(1.013)
Directional network exposure, N_{it}^{Δ}	-0.225	-0.220	-0.225	0.012
	(0.253)	(0.252)	(0.270)	(0.280)
$N_{it}^{\Delta} \times \pi_i$	-0.077	-0.071	-0.077	-0.821
	(0.557)	(0.559)	(0.574)	(0.634)
Regional drought, D_{it}^{reg}	0.604	0.583	0.604	0.749
	(0.505)	(0.507)	(0.610)	(0.519)
$\pi_i \times (t - \bar{t})$	-1.162^{***}	-1.165^{***}	-1.162^{***}	-1.036^{***}
	(0.192)	(0.192)	(0.196)	(0.185)
Observations	15,345	15,190	15,345	13,810
Counties	3,069	3,038	3,069	2,762
R^2	0.953	0.953	0.953	0.953
County fixed effects	Yes	Yes	Yes	Yes
Election fixed effects	Yes	Yes	Yes	Yes
Conley spatial cutoff	500 km	500 km	800 km	500 km
Joint Wald test of shock terms, W/q	6.177	6.142	5.564	7.772
p-value	< 0.001	< 0.001	< 0.001	< 0.001

Notes: All columns estimate the all-channels benchmark specification with county and election fixed effects. Column (2) excludes the top 1% of counties by total external food flows, column (3) uses an 800 km rather than 500 km Conley cutoff, and column (4) excludes the top HHI decile. Network weights are not recomputed after sample exclusions. Conley standard errors are reported in parentheses. The Wald test evaluates the joint significance of the drought-exposure terms. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A separate concern concerns spatial inference. Food-flow connections extend beyond immediately neighboring counties, and residual dependence may therefore persist over distances greater than the 500 km cutoff used in the benchmark Conley covariance estimator. Column (3) consequently increases the spatial cutoff from 500 to 800 km while leaving the estimation sample and regression specification unchanged. Because this

exercise changes only the covariance estimator, the point estimates are identical to those in the benchmark column.

The estimates are stable across these exercises. Excluding the largest network hubs leaves both central partisan gradients almost unchanged: the local-drought interaction moves from 2.174 in the benchmark to 2.158, while the mean-network interaction changes from 2.596 to 2.637. Both remain statistically significant at the 1 percent level. The similarity of these estimates indicates that the benchmark relationships are not driven by a small number of counties with exceptionally large total food flows.

The results are also robust to excluding counties with highly concentrated partner weights. The local-drought interaction remains close to its benchmark value, at 2.148, while the mean-network interaction increases to 3.126 and remains statistically significant at the 1 percent level. Thus, the positive partisan gradient associated with mean network drought does not depend on counties whose exposure is dominated by a particularly small set of partners. The directional network coefficient and its interaction with historical Democratic support remain statistically insignificant, providing no evidence that the supplier–destination composition of network exposure becomes systematically important once highly concentrated networks are removed.

Finally, increasing the Conley spatial cutoff from 500 to 800 km has little effect on inference. The standard errors change only modestly, and the local and mean-network partisan gradients remain statistically significant at the 1 percent level. The drought-exposure terms are also jointly significant in all four specifications ($p < 0.001$). Taken together, these exercises indicate that the benchmark results are not driven by network hubs, unusually concentrated bilateral weights, or the particular spatial bandwidth used for Conley inference.

6.3 Sensitivity to the 2012 and 2016 elections

Because the panel contains only five presidential elections, an additional concern is that the estimated partisan gradients may be disproportionately influenced by a particular election. We therefore conduct two targeted leave-one-election-out exercises, motivated by distinct sources of election-specific variation.

The first excludes 2012. As discussed in Section 3.2, pre-election drought was substantially more widespread and severe in 2012 than in any other election in the sample. More than 80 percent of counties experienced positive drought exposure, and mean standardized drought intensity was more than twice the full-sample average. Removing 2012 therefore provides a demanding test of whether the benchmark relationships depend on the unusually large cross-county variation in drought associated with that election.

The second exercise excludes 2016. The motivation is political rather than climatic. Donald Trump's first presidential candidacy coincided with a marked reconfiguration of

U.S. electoral geography, including particularly strong Democratic losses along the rural side of the urban–rural continuum (Scala and Johnson, 2017). Other studies document distinctive electoral divisions associated with education, perceived status threat, racial attitudes, sexism, and immigration (Mutz, 2018; Schaffner et al., 2018; Hooghe and Dassonneville, 2018). Although election fixed effects and the differential partisan trend absorb common election-specific shocks and systematic partisan change, 2016 may still contain distinctive cross-sectional variation correlated with the geography of drought exposure. Excluding 2016 therefore tests whether the estimated drought gradients are sensitive to this unusually distinctive electoral episode.

For each exercise, we re-estimate the benchmark specification after removing the corresponding election, without otherwise modifying the empirical model. All specifications include county and election fixed effects, regional drought D_{it}^{reg}, and the differential partisan trend $\pi_i(t - \bar{t})$. Conley standard errors are calculated using the same spatial and temporal structure as in the benchmark specification. To preserve comparability across samples, all drought-exposure variables retain the normalization calculated from the full benchmark sample. Each restricted sample therefore contains four elections and 12,276 county–election observations.

Table 11 shows that the two exclusions have different implications for precision but do not reverse the sign of the estimated interactions with historical Democratic support.

When 2012 is excluded, the interactions of both local drought and mean network exposure with historical Democratic support remain positive, at 1.743 and 1.755, respectively. Neither is individually significant at conventional levels, however, and the joint Wald test rejects the null only at the 10 percent level ($p = 0.094$). The reduction in precision is substantial, but it occurs after removing one of only five elections and, in particular, the election providing the greatest cross-county variation in pre-election drought. The persistence of positive partisan gradients therefore suggests that 2012 contributes importantly to statistical precision without determining the direction of the estimated relationships.

The results are stronger when 2016 is excluded. The interaction between mean network exposure and historical Democratic support increases from 2.596 in the benchmark to 3.105 and remains statistically significant at the 1 percent level. The corresponding local-drought interaction remains positive, at 1.379, although it is no longer individually significant. The drought-exposure terms are nevertheless jointly highly significant ($p < 0.001$), while neither directional network exposure nor its interaction with historical Democratic support is statistically significant. The stability of the mean-network interaction therefore indicates that this relationship does not appear to be driven by the distinctive electoral geography associated with Trump's first presidential candidacy. The local relationship is less precisely estimated once 2016 is removed, but its partisan gradient retains the sign predicted by the model.

Table 11: Sensitivity to the exclusion of the 2012 and 2016 elections

	Benchmark	Excluding 2012	Excluding 2016
Local drought, D_{it}	-0.879^{***}	-0.823^{**}	-0.642^{*}
	(0.317)	(0.419)	(0.359)
$D_{it} \times \pi_i$	2.174^{***}	1.743	1.379
	(0.791)	(1.093)	(0.898)
Mean network exposure, N_{it}^{m}	-1.135^{***}	-0.717	-1.345^{***}
	(0.396)	(0.488)	(0.405)
$N_{it}^{m} \times \pi_i$	2.596^{***}	1.755	3.105^{***}
	(0.885)	(1.186)	(0.896)
Directional network exposure, N_{it}^{Δ}	-0.225	0.332	-0.043
	(0.253)	(0.375)	(0.291)
$N_{it}^{\Delta} \times \pi_i$	-0.077	-0.467	-0.353
	(0.557)	(0.858)	(0.651)
Regional drought, D_{it}^{reg}	0.604	-0.621	0.731
	(0.505)	(0.681)	(0.652)
$\pi_i \times (t - \bar{t})$	-1.162^{***}	-1.172^{***}	-1.166^{***}
	(0.192)	(0.231)	(0.181)
Observations	15,345	12,276	12,276
R^2	0.953	0.956	0.948
Joint Wald test, p-value	< 0.001	0.094	< 0.001

Notes: The dependent variable is the Democratic two-party presidential vote share. All columns report the all-channels specification with county and election fixed effects, regional drought, and the differential partisan trend $\pi_i(t - \bar{t})$. Exposure variables retain the full-sample normalization. Conley standard errors are reported in parentheses. The Wald test jointly evaluates the level and partisan-interaction coefficients for all three drought-exposure channels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The implied marginal effects reinforce this qualitative pattern. In both restricted samples, local and mean network drought are associated with negative estimated effects at low levels of historical Democratic support, effects close to zero near the center of the partisan distribution, and positive effects at high levels. When 2012 is excluded, the individual marginal effects are less precisely estimated, although the combined local-and-network effect remains negative and significant at $\pi_i = 0.25$ and positive at the 10 percent level at $\pi_i = 0.75$. When 2016 is excluded, the corresponding mean-network and combined effects remain precisely estimated at both ends of the partisan distribution. Full marginal-effect estimates are reported in Appendix Table C.1.

Because excluding 2012 removes one of only five presidential elections as well as the election with the strongest cross-county variation in drought, we complement the leave-one-election-out exercise with a pooled heterogeneity specification that retains all five elections. Specifically, for each exposure channel, we allow both the level coefficient and its interaction with historical Democratic support to differ in 2012. The resulting tests distinguish between a 2012-specific difference in the partisan gradient and a difference

in the entire slope function. Appendix Table C.2 summarizes the implied gradients and corresponding Wald tests, while Appendix Table C.3 reports the complete coefficient estimates.

The estimated partisan gradient for local drought is 1.907 outside 2012 and 2.367 in 2012, and the difference is not statistically significant ($p = 0.759$). For mean network exposure, the corresponding gradients are 2.054 and 2.021, with an estimated difference close to zero ($p = 0.987$). A joint test of equality of the local and mean-network partisan gradients also fails to reject the null ($p = 0.948$). More broadly, a joint test allowing both the level and partisan-gradient components of the local and mean-network relationships to differ in 2012 does not reject equality ($p = 0.461$). We therefore find no evidence that the two central relationships identified in the benchmark differ systematically in 2012 from those observed across the remaining elections.

The directional network component provides one qualification. Its partisan gradient does not differ significantly between 2012 and the other elections ($p = 0.381$), but the joint test that both its level coefficient and partisan gradient are unchanged rejects strongly ($p < 0.001$). This indicates some 2012-specific variation in the directional component, although not specifically in its partisan gradient. Because directional exposure is not statistically significant in the benchmark estimates, this qualification does not alter the evidence concerning the two central relationships, but it suggests that the supplier–destination composition of network exposure differed in 2012.

Taken together, these exercises provide no evidence that the qualitative partisan pattern documented in the benchmark is specific either to the exceptional 2012 drought or to the distinctive political environment of the 2016 election. At the same time, they underscore an important limitation of the analysis: with only five presidential elections, the precision of some estimates—particularly the local-drought gradient—is sensitive to the temporal composition of the sample.

7. Alternative Interpretations and Electoral Channels

7.1 Incumbent-party heterogeneity and retrospective punishment

The main analysis shows that the relationships between drought exposure and Democratic vote share vary with counties' historical Democratic support. We next ask whether these relationships also differ with the party holding the presidency in the direction predicted by simple retrospective punishment. Under that interpretation, adverse drought exposure should be more unfavorable to the Democratic vote share when a Democratic administration is in office and more favorable when a Republican administration is in office.

To examine this prediction while preserving the benchmark interaction with historical Democratic support, we define a signed indicator for the party holding the presidency immediately before election t:

$$\rho_t = \begin{cases} +1, & \text{if the incumbent presidential party is Democratic,} \\ -1, & \text{if the incumbent presidential party is Republican.} \end{cases} \tag{26}$$

Accordingly, $\rho_t = -1$ in 2008 and 2020 and $\rho_t = +1$ in 2012, 2016, and 2024. The coding refers to the party controlling the presidency before the election, irrespective of whether the incumbent president is personally seeking reelection.

We estimate the following extension of the benchmark specification:

$$y_{it} = \alpha_i + \lambda_t + \sum_{k \in K} X_{it}^k \left[\beta_{0k} + \beta_{1k}\pi_i + \gamma_k \rho_t\right] + \theta D_{it}^{reg} + \eta \pi_i (t - \bar{t}) + \varepsilon_{it}, \tag{27}$$

where $K = \{D, N^m, N^\Delta\}$. Election fixed effects absorb the main effect of ρ_t. The coefficient γ_k therefore captures how the relationship between exposure channel k and Democratic vote share shifts between Democratic and Republican administrations, holding historical Democratic support fixed. Because the specification does not include $X_{it}^k \pi_i \rho_t$, this shift is constrained to be constant across values of π_i.

For a given π_i, the marginal relationships under Democratic and Republican administrations are:

$$ME_k^{Dem}(\pi_i) = \beta_{0k} + \beta_{1k}\pi_i + \gamma_k \tag{28}$$

$$ME_k^{Rep}(\pi_i) = \beta_{0k} + \beta_{1k}\pi_i - \gamma_k \tag{29}$$

so that,

$$ME_k^{Dem}(\pi_i) - ME_k^{Rep}(\pi_i) = 2\gamma_k \tag{30}$$

Thus, γ_k is one half of the estimated difference between the two presidential-party regimes rather than a marginal effect of drought itself. For local drought and mean network exposure, simple retrospective punishment predicts:

$$\gamma_D < 0 \qquad \text{and} \qquad \gamma_{N^m} < 0,$$

because a given increase in drought exposure should reduce Democratic support more under a Democratic administration than under a Republican one. No corresponding sign restriction is imposed on γ_{N^Δ} because the directional measure captures the supplier–destination composition of network exposure rather than overall drought intensity.

Table 12 reports the estimates.

Table 12: Drought-exposure relationships by the party holding the presidency

Drought-exposure channels	Local only D only	Local and network D, N^m	Local and asymmetric network D, N^Δ	All channels D, N^m, N^Δ
Local drought exposure, D_{it}	-1.643^{***}	-0.921^{***}	-1.662^{***}	-0.960^{***}
	(0.311)	(0.314)	(0.308)	(0.311)
$D_{it} \times \pi_i$	3.647^{***}	2.042^{***}	3.669^{***}	2.091^{***}
	(0.741)	(0.791)	(0.734)	(0.784)
$D_{it} \times \rho_t$ (half-difference between regimes)	0.348^{**}	0.320^{***}	0.343^{**}	0.311^{***}
	(0.146)	(0.110)	(0.146)	(0.111)
Mean network exposure, N^m_{it}	–	-1.258^{***}	–	-1.220^{***}
	–	(0.422)	–	(0.409)
$N^m_{it} \times \pi_i$	–	2.757^{***}	–	2.738^{***}
	–	(0.922)	–	(0.893)
$N^m_{it} \times \rho_t$ (half-difference between regimes)	–	0.063	–	0.071
	–	(0.148)	–	(0.149)
Asymmetric network exposure, N^Δ_{it}	–	–	-0.384	-0.178
	–	–	(0.263)	(0.244)
$N^\Delta_{it} \times \pi_i$	–	–	0.420	-0.054
	–	–	(0.599)	(0.553)
$N^\Delta_{it} \times \rho_t$ (half-difference between regimes)	–	–	-0.075	-0.074
	–	–	(0.064)	(0.065)
Regional control, D^{reg}_{it}	0.521	0.554	0.516	0.524
	(0.479)	(0.509)	(0.477)	(0.508)
$\pi_i \times (t - \bar{t})$	-1.201^{***}	-1.175^{***}	-1.199^{***}	-1.175^{***}
	(0.188)	(0.190)	(0.188)	(0.190)
Observations	15 345	15 345	15 345	15 345
R^2	0.953	0.953	0.953	0.953
Joint Wald test of incumbent-party interactions, W/q	5.694	4.211	4.303	3.463
p-value	0.017	0.015	0.014	0.016

Note: The dependent variable is the Democratic two-party vote share. All specifications include county and election fixed effects. $\rho_t = 1$ under a Democratic administration and -1 under a Republican administration, so the coefficient on $X^k_{it} \times \rho_t$ equals one half of the difference between the two regimes. Conley standard errors are reported in parentheses. The Wald test jointly evaluates the included incumbent-party interactions. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

In the all-channels specification, the incumbent-party interactions are jointly significant ($W/q = 3.463$, $p = 0.016$), indicating that the estimated drought relationships differ across the two presidential-party regimes. The sign of these differences, however, does not conform to the prediction of simple retrospective punishment.

For local drought, $\widehat{\gamma}_D = 0.311$ and is statistically significant at the 1 percent level. The implied difference between Democratic and Republican administrations is therefore $2\widehat{\gamma}_D = 0.622$ percentage points. Its positive sign means that, holding π_i fixed, the estimated local-drought relationship is more favorable to the Democratic vote share under a Democratic administration than under a Republican administration—the opposite of the sign predicted by retrospective punishment. The estimate is similar across the four exposure specifications, ranging from 0.311 to 0.348.

For mean network exposure, the corresponding regime difference is much smaller: $2\widehat{\gamma}_{N^m} = 0.143$, with $p = 0.632$. There is therefore little evidence that the relationship between mean network drought and Democratic vote share differs systematically with the party controlling the presidency. For a simultaneous one-unit increase in local drought and mean network exposure, the estimated regime difference is $2(\widehat{\gamma}_D + \widehat{\gamma}_{N^m}) = 0.764$ percentage points. This difference is again positive, rather than negative as predicted by retrospective punishment, but is only marginally significant ($p = 0.054$).

Table 13 reports the corresponding marginal effects under Democratic and Republican administrations at selected values of historical Democratic support.

Table 13: Marginal effects of drought by historical Democratic vote share and the party holding the presidency

Drought-exposure channel	Incumbent-party regime	Historical Democratic support, π_i				
		$0.00^{\dagger}$	0.25	0.50	0.75	$1.00^{\dagger}$
Local drought, D_{it}	Democratic incumbent	-0.649^{**}	-0.127	0.396^{***}	0.919^{***}	1.442^{***}
		(0.326)	(0.158)	(0.144)	(0.305)	(0.493)
	Republican incumbent	-1.271^{***}	-0.748^{***}	-0.226	0.297	0.820
		(0.334)	(0.210)	(0.232)	(0.374)	(0.551)
	Democratic − Republican	0.622^{***}	0.622^{***}	0.622^{***}	0.622^{***}	0.622^{***}
		(0.221)	(0.221)	(0.221)	(0.221)	(0.221)
Mean network exposure, N_{it}^m	Democratic incumbent	-1.148^{***}	-0.464^{**}	0.221	0.906^{***}	1.590^{***}
		(0.405)	(0.209)	(0.152)	(0.319)	(0.529)
	Republican incumbent	-1.291^{***}	-0.606^{**}	0.078	0.763^{**}	1.448^{**}
		(0.464)	(0.305)	(0.266)	(0.385)	(0.570)
	Democratic − Republican	0.143	0.143	0.143	0.143	0.143
		(0.298)	(0.298)	(0.298)	(0.298)	(0.298)
Combined effect D_{it},N_{it}^m	Democratic incumbent	-1.798^{***}	-0.590^{**}	0.617^{***}	1.824^{***}	3.032^{***}
		(0.421)	(0.249)	(0.211)	(0.354)	(0.552)
	Republican incumbent	-2.562^{***}	-1.355^{***}	-0.147	1.060^{**}	2.267^{***}
		(0.491)	(0.384)	(0.389)	(0.504)	(0.675)
	Democratic − Republican	0.764^{*}	0.764^{*}	0.764^{*}	0.764^{*}	0.764^{*}
		(0.396)	(0.396)	(0.396)	(0.396)	(0.396)

Note: Marginal effects are from the all-channels incumbent-party specification. The reported difference is the effect under a Democratic administration minus that under a Republican administration and is constant across π_i because $X_{it}^k \pi_i \rho_t$ is not included. Conley standard errors are reported in parentheses. Values marked with † lie outside the observed range of π_i. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Under both presidential-party regimes, the estimated relationships become more favorable to the Democratic vote share as historical Democratic support increases, reproducing the sign-changing pattern found in the benchmark analysis. For local drought, the marginal effect is uniformly 0.622 percentage points more positive under a Democratic administration, whereas the corresponding difference for mean network exposure is only 0.143 percentage points and is statistically insignificant. The combined local-and-network relationship displays the same pattern, with an estimated regime difference of 0.764 percentage points. Thus, presidential-party control shifts the level of some estimated drought relationships but does not alter their systematic variation with historical Democratic

support.

These comparisons should be interpreted cautiously. Presidential-party control varies only across the five elections in the sample—two held under Republican administrations and three under Democratic administrations—so the specification cannot isolate presidential-party control from other election-specific conditions. The exercise is therefore descriptive: it asks whether the differences between the two sets of elections have the sign implied by retrospective punishment rather than identifying a causal incumbent-party effect.

Taken together, the estimates do not display the regime pattern predicted by simple retrospective punishment. The local-drought difference is estimated precisely but has the opposite sign, while the mean-network difference is small and statistically insignificant. The more systematic feature of the results remains the variation in drought relationships with historical Democratic support documented in the benchmark analysis.

7.2 Turnout versus partisan composition

The preceding results use the Democratic two-party vote share as the dependent variable. To assess whether these results reflect changes in the partisan composition of votes or in electoral participation, we re-estimate the same specification using three outcomes measured as shares of the citizen voting-age population (CVAP): total turnout, Democratic votes, and Republican votes. CVAP data are obtained from the U.S. Census Bureau's American Community Survey. Table 14 reports the baseline estimates, while Table 15 reports estimates from the specification that includes the differential nonmetropolitan partisan trend, $\pi_i R_i \left(t - \bar{t}\right)$.[10]

Table 14 provides little evidence that the drought relationships operate through aggregate electoral participation. For turnout as a share of CVAP, the coefficients on local drought, mean network exposure, and their interactions with historical Democratic vote share are small relative to their standard errors. The joint tests also fail to reject the null that all included drought terms are jointly equal to zero. The turnout-test p-values are 0.834 in the local-drought model and 0.504 in the full model. Although the directional network terms are individually significant at the 10 percent level in the full specification, the drought terms are not jointly significant. Taken together, these estimates provide no consistent evidence that drought exposure is associated with aggregate turnout.

By contrast, Democratic votes as a share of CVAP show a pattern similar to that obtained for the Democratic two-party vote share. In the full baseline specification, the coefficients on local drought and mean network exposure are negative, at -0.730 and -0.764, respectively, while their interactions with historical Democratic vote share are positive, at 1.400 and 1.599. All four coefficients are statistically significant at the 1

[10]The CVAP-based sample contains 15,337 county-election observations because CVAP is unavailable for eight observations in the benchmark electoral sample.

percent level. The directional network exposure measure and its interaction are not statistically significant. The drought-related terms are also jointly significant in both Democratic-vote specifications ($p < 0.001$).

Table 14: Drought exposure, turnout and party votes as shares of CVAP

	Turnout / CVAP		Dem. votes / CVAP		Rep. votes / CVAP	
Drought-exposure channels	D only	D, N^m, N^Δ	D only	D, N^m, N^Δ	D only	D, N^m, N^Δ
Local drought exposure, D_{it}	−0.307	−0.240	−1.153***	−0.730***	0.554	0.449
	(1.414)	(1.217)	(0.228)	(0.220)	(0.365)	(0.321)
$D_{it} \times \pi_i$	0.246	0.066	2.325***	1.400***	−1.739**	−1.422**
	(2.505)	(2.217)	(0.485)	(0.493)	(0.750)	(0.712)
Network exposure, N^m_{it}	–	−0.078	–	−0.764***	–	0.187
	–	(1.251)	–	(0.244)	–	(0.448)
$N^m_{it} \times \pi_i$	–	0.158	–	1.599***	–	−0.631
	–	(2.453)	–	(0.534)	–	(0.923)
Asymmetric network exposure, N^Δ_{it}	–	−1.657*	–	−0.134	–	−0.325
	–	(0.924)	–	(0.185)	–	(0.459)
$N^\Delta_{it} \times \pi_i$	–	3.321*	–	0.058	–	1.086
	–	(1.853)	–	(0.405)	–	(0.930)
Regional control, D^{reg}	1.897**	1.916*	1.140***	1.190***	0.332	0.394
	(0.913)	(1.012)	(0.343)	(0.364)	(0.340)	(0.360)
$\pi_i \times (t - \bar{t})$	−2.752***	−2.739***	−0.643***	−0.626***	−0.016	−0.019
	(0.955)	(0.957)	(0.095)	(0.096)	(0.166)	(0.167)
Observations	15 337	15 337	15 337	15 337	15 337	15 337
R^2	0.607	0.607	0.950	0.950	0.910	0.910
Joint Wald test of shock terms, W/q	0.181	0.886	13.079	5.515	4.898	2.543
p-value	$p = 0.834$	$p = 0.504$	$p < 0.001$	$p < 0.001$	$p = 0.007$	$p = 0.018$

Note: The dependent variables are turnout, Democratic votes, and Republican votes, each as a share of CVAP. Conley standard errors are reported in parentheses. The joint Wald test evaluates whether the drought-exposure measures and their interactions with historical Democratic support included in each column are jointly equal to zero; W/q denotes the normalized Wald statistic. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Republican votes as a share of CVAP display a less systematic pattern. In Table 14, the interaction between local drought and historical Democratic vote share is negative and statistically significant in both the local-drought and full specifications, at −1.739 and −1.422, respectively. Thus, the relationship between local drought and Republican votes becomes more negative as historical Democratic support increases. However, neither the coefficient on mean network exposure nor its interaction with historical Democratic support is statistically significant.

Table 15 shows that allowing for differential nonmetropolitan partisan trends leaves the Democratic-vote results largely unchanged. In the full specification, the coefficients on local drought and mean network exposure remain negative, while their interactions with historical Democratic support remain positive and statistically significant. The

drought-related terms remain jointly significant for Democratic votes ($p < 0.001$), and the directional network terms remain statistically insignificant.

Table 15: Drought exposure, turnout and party votes as shares of CVAP, with a differential nonmetropolitan trend

	Turnout / CVAP		Dem. votes / CVAP		Rep. votes / CVAP	
Drought-exposure channels	D only	D, N^m, N^Δ	D only	D, N^m, N^Δ	D only	D, N^m, N^Δ
Local drought exposure, D_{it}	−0.324	−0.259	−1.040***	−0.602***	0.434	0.314
	(1.407)	(1.212)	(0.211)	(0.202)	(0.359)	(0.309)
$D_{it} \times \pi_i$	0.286	0.112	2.062***	1.090***	−1.460**	−1.093
	(2.490)	(2.206)	(0.455)	(0.460)	(0.730)	(0.670)
Network exposure, N_{it}^m	–	−0.077	–	−0.774***	–	0.198
	–	(1.252)	–	(0.239)	–	(0.445)
$N_{it}^m \times \pi_i$	–	0.149	–	1.659***	–	−0.694
	–	(2.458)	–	(0.533)	–	(0.923)
Asymmetric network	–	−1.642*	–	−0.229	–	−0.224
exposure, N_{it}^Δ	–	(0.920)	–	(0.177)	–	(0.459)
$N_{it}^\Delta \times \pi_i$	–	3.280*	–	0.335	–	0.792
	–	(1.841)	–	(0.390)	–	(0.937)
Regional control, D^{reg}	1.904**	1.924*	1.096***	1.140***	0.379	0.448
	(0.912)	(1.010)	(0.320)	(0.339)	(0.326)	(0.344)
$\pi_i \times (t - \bar{t})$	−2.803***	−2.789***	−0.313***	−0.295***	−0.367**	−0.370**
	(0.970)	(0.973)	(0.098)	(0.099)	(0.164)	(0.168)
$\pi_i \times R_i \times (t - \bar{t})$	0.089	0.087	−0.585***	−0.585***	0.622***	0.621***
	(0.095)	(0.095)	(0.043)	(0.043)	(0.043)	(0.043)
Observations	15 337	15 337	15 337	15 337	15 337	15 337
R^2	0.607	0.607	0.954	0.954	0.914	0.914
Joint Wald test of shock terms, W/q	0.172	0.868	12.285	5.006	4.367	1.979
p-value	$p = 0.842$	$p = 0.517$	$p < 0.001$	$p < 0.001$	$p = 0.013$	$p = 0.065$

Note: The dependent variables are turnout, Democratic votes, and Republican votes, each as a share of CVAP. All specifications additionally include the differential nonmetropolitan partisan trend $\pi_i R_i (t - \bar{t})$. Conley standard errors are reported in parentheses. The joint Wald test evaluates whether the drought-exposure measures and their interactions with historical Democratic support included in each column are jointly equal to zero; W/q denotes the normalized Wald statistic. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The turnout results are also essentially unchanged: the joint tests remain insignificant in both the drought-only and full specifications ($p = 0.842$ and $p = 0.517$, respectively).

The Republican-vote results are more sensitive to this additional trend. The local-drought interaction remains negative and statistically significant in the drought-only specification but declines to −1.093 and is no longer statistically significant in the full model. The joint test in the full Republican-vote specification rejects at the 10 percent level but not at the 5 percent level ($p = 0.065$). Thus, the evidence for Republican votes is weaker and less stable than the corresponding evidence for Democratic votes.

Taken together, the CVAP results point more strongly to changes in the partisan

composition of votes than to a common change in aggregate turnout. The sign-changing relationship observed for the Democratic two-party vote share is reproduced more consistently for Democratic votes as a share of CVAP, whereas the Republican-vote results weaken after accounting for differential nonmetropolitan partisan trends. This asymmetry does not contradict the two-party vote-share results. Mechanically, the Democratic two-party vote share can increase because Democratic votes rise, Republican votes fall, or both. The CVAP results indicate that the first component exhibits the estimated drought-related pattern more consistently than the second.

8. Conclusion

This paper examines whether drought is associated with electoral outcomes only where it occurs or also in counties connected to drought-affected locations through food-flow networks. We combine county-level presidential election returns from 2008 to 2024 with pre-election drought measures and bilateral food-flow weights based on estimated 2007 flows. The empirical framework distinguishes a county's own drought from drought among its food-flow partners and allows both relationships to vary with historical Democratic vote share. The analysis therefore extends the political economy of climate shocks beyond drought occurring within a county's own boundaries.

The results show a consistent pattern of partisan heterogeneity. Local drought is associated with a lower Democratic two-party vote share in counties with low historical Democratic vote shares and a higher Democratic vote share in counties with high historical Democratic vote shares. Mean drought exposure among food-flow partners exhibits a similar sign-changing relationship. The estimated zero-crossing points for local and network exposure are close and lie near the center of the distribution of historical Democratic support. The implied relationships are therefore small for counties near the center of that distribution and become increasingly different in sign and magnitude toward its two ends. The main result is thus one of partisan heterogeneity rather than a uniform electoral shift toward either party. By contrast, the directional network component—which distinguishes exposure through suppliers from exposure through destination markets—is generally small and statistically not significant. The evidence is therefore stronger for overall drought exposure among connected counties than for differences based on the direction of food flows.

The network relationship is also present in county-election observations with no measured local drought. Counties whose local drought measure equals zero but that have high network exposure show a sign-changing relationship across historical Democratic vote shares similar to that found in the full sample. Moreover, the estimated relationship is larger for high network exposure than for low exposure. These results show that the network pattern is not confined to county-election observations that simultaneously

experience measured local drought. They provide additional evidence that drought among economically connected counties is associated with electoral outcomes even when the county's six-month average local-drought measure equals zero.

Allowing historical Democratic support to follow a differential trend in nonmetropolitan counties reduces the magnitude of the local-drought interaction but leaves the mean-network interaction largely unchanged. The sign-changing pattern and the proximity of the local and network zero-crossing points are also preserved under this augmented specification. The main results remain stable under further robustness and sensitivity exercises, including alternative election-specific fixed effects, restrictions on network hubs and concentrated weights, a wider spatial cutoff, and changes in the election sample. Across these exercises, the mean-network gradient is generally more stable than the local-drought gradient, whose magnitude and precision are more sensitive to the limited number of presidential elections in the sample.

The decomposition of electoral outcomes helps clarify whether these results reflect turnout or party-specific vote totals. Drought-related coefficients are not jointly significant for aggregate turnout as a share of the citizen voting-age population. The main sign-changing relationship is more consistently observed in Democratic votes as a share of the citizen voting-age population, while the results for Republican votes are less stable after accounting for differential partisan trends in nonmetropolitan counties. Taken together, these estimates are more consistent with changes in party-specific electoral support than with a common change in aggregate turnout. They do not, however, establish whether this compositional change reflects individual vote switching, differential mobilization of partisan constituencies, or a combination of the two.

The comparison across presidential-party regimes provides little support for a simple retrospective-punishment interpretation. A retrospective account would predict drought to be more unfavorable to the Democratic vote share during Democratic administrations, but the estimated differences do not have this sign.

These findings contribute to our understanding of how climate-related disturbances become politically consequential. The estimates are consistent with partisan differences in the electoral response to drought-related exposure. Rather than producing a uniform shift toward one party or systematic punishment of the party holding the presidency, drought exposure is associated with electoral changes in opposite directions across counties with different partisan histories. This pattern is consistent with the partisan-attribution framework developed in the paper, but it does not uniquely identify belief updating as the underlying mechanism. Other processes that vary with historical partisan alignment, including economic exposure, mobilization, or differences in information and political cues, could contribute to the same reduced-form pattern.

The similarity between the local and network relationships suggests that the political relevance of climate exposure may extend beyond the places where the physical event

occurs. Food-flow connections provide one dimension along which drought outside a county may become associated with electoral outcomes within it, including when the measured local-drought index equals zero.

The results also motivate a broader view of geographic exposure to climate risk. If economic connections transmit some of the consequences of drought, measures based only on local physical conditions may provide an incomplete picture of the places potentially exposed to those consequences. Incorporating information on supply-chain connections into climate-risk assessment may therefore help identify forms of indirect exposure that are not captured by local weather measures alone.

The partisan heterogeneity of the estimates also points to a possible role for differences in how drought-related information is received and interpreted across communities. The present analysis cannot determine whether these differences arise from partisan information processing, heterogeneous economic effects, mobilization, or other mechanisms. Understanding these channels is therefore an important direction for future work.

Future research could connect the electoral estimates more directly to economic mechanisms by combining food-flow exposure with county-level data on agricultural production, employment, income, retail food prices, transportation disruptions, and government assistance. Individual-level surveys, combined where possible with validated turnout records, could help distinguish changes in vote choice from changes in electoral participation and examine whether residents recognize economic consequences associated with drought in counties connected through food flows.

Overall, the evidence indicates that the electoral relevance of drought is neither strictly local, uniform across partisan contexts, nor well characterized by a simple incumbent-punishment response. Drought exposure is associated with electoral changes that vary systematically with a county's historical Democratic vote share, and similar relationships appear when drought occurs among the county's food-flow partners rather than within the county itself. The central empirical finding is therefore a sign-changing partisan relationship that extends from local drought to drought exposure through pre-sample food-flow connections. Climate shocks may be politically relevant across a wider economic geography than their physical footprint alone suggests.

References

Aadland, D. (2004). Cattle cycles, heterogeneous expectations and the age distribution of capital. *Journal of Economic Dynamics and Control*, 28(10):1977–2002.

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The Network Origins of Aggregate Fluctuations. *Econometrica*, 80(5):1977–2016.

Achen, C. H. and Bartels, L. M. (2004). Blind Retrospection: Electoral Responses to Drought, Flu, and Shark Attacks. Working Paper 2004/199, Center for Advanced Study in the Social Sciences, Juan March Institute, Madrid. Later published as Chapter 5 of *Democracy for Realists*, Princeton University Press, 2016.

Achen, C. H. and Bartels, L. M. (2016). *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton University Press, Princeton, NJ.

Adão, R., Kolesár, M., and Morales, E. (2019). Shift-Share Designs: Theory and Inference. *Quarterly Journal of Economics*, 134(4):1949–2010.

Amlani, S. and Algara, C. (2021). Partisanship and nationalization in American elections: Evidence from presidential, senatorial, and gubernatorial elections in the U.S. counties, 1872–2020. *Electoral Studies*, 73:102387.

Autor, D., Dorn, D., Hanson, G. H., and Majlesi, K. (2020). Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure. *American Economic Review*, 110(10):3139–3183.

Barrot, J.-N. and Sauvagnat, J. (2016). Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks. *Quarterly Journal of Economics*, 131(3):1543–1592.

Basara, J. B., Christian, J. I., Wakefield, R. A., Otkin, J. A., Hunt, E. H., and Brown, D. P. (2019). The evolution, propagation, and spread of flash drought in the Central United States during 2012. *Environmental Research Letters*, 14(8):084025.

Bechtel, M. M. and Hainmueller, J. (2011). How Lasting Is Voter Gratitude? An Analysis of the Short- and Long-Term Electoral Returns to Beneficial Policy. *American Journal of Political Science*, 55(4):852–868.

Bisgaard, M. (2015). Bias Will Find a Way: Economic Perceptions, Attributions of Blame, and Partisan-Motivated Reasoning during Crisis. *Journal of Politics*, 77(3):849–860.

Bisgaard, M. (2019). How Getting the Facts Right Can Fuel Partisan-Motivated Reasoning. *American Journal of Political Science*, 63(4):824–839.

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake. *Review of Economics and Statistics*, 101(1):60–75.

Bolsen, T., Druckman, J. N., and Cook, F. L. (2014). The Influence of Partisan Motivated Reasoning on Public Opinion. *Political Behavior*, 36(2):235–262.

Bonomi, G., Gennaioli, N., and Tabellini, G. (2021). Identity, Beliefs, and Political Conflict. *Quarterly Journal of Economics*, 136(4):2371–2411.

Brown, T. E. and Mettler, S. (2024). Sequential Polarization: The Development of the Rural–Urban Political Divide, 1976–2020. *Perspectives on Politics*, 22(3):630–658.

Bruckner, M., Wood, R., Moran, D., Kuschnig, N., Wieland, H., Maus, V., and Börner, J. (2019). FABIO—The Construction of the Food and Agriculture Biomass Input–Output Model. *Environmental Science & Technology*, 53(19):11302–11312.

Bullock, J. G., Gerber, A. S., Hill, S. J., and Huber, G. A. (2015). Partisan Bias in Factual Beliefs about Politics. *Quarterly Journal of Political Science*, 10(4):519–578.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply Chain Disruptions: Evidence from the Great East Japan Earthquake. *Quarterly Journal of Economics*, 136(2):1255–1321.

Chavas, J.-P. (2000). On information and market dynamics: The case of the U.S. beef market. *Journal of Economic Dynamics and Control*, 24(5–7):833–853.

Cole, S., Healy, A., and Werker, E. (2012). Do voters demand responsive governments? evidence from Indian disaster relief. *Journal of Development Economics*, 97(2):167–181.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics*, 92(1):1–45.

Deryugina, T. (2013). How do people update? the effects of local weather fluctuations on beliefs about global warming. *Climatic Change*, 118(2):397–416.

Djourelova, M., Durante, R., Motte, E., and Patacchini, E. (2023). Experience, Narratives, and Climate Change Beliefs. *SSRN Electronic Journal*.

Druckman, J. N. and McGrath, M. C. (2019). The evidence for motivated reasoning in climate change preference formation. *Nature Climate Change*, 9(2):111–119.

Egan, P. J. and Mullin, M. (2012). Turning Personal Experience into Political Attitudes: The Effect of Local Weather on Americans' Perceptions about Global Warming. *Journal of Politics*, 74(3):796–809.

Egan, P. J. and Mullin, M. (2017). Climate Change: US Public Opinion. *Annual Review of Political Science*, 20:209–227.

Fair, R. C. (1978). The Effect of Economic Events on Votes for President. *Review of Economics and Statistics*, 60(2):159–173.

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

Gasper, J. T. and Reeves, A. (2011). Make It Rain? Retrospection and the Attentive Electorate in the Context of Natural Disasters. *American Journal of Political Science*, 55(2):340–355.

Hart, P. S. and Nisbet, E. C. (2012). Boomerang Effects in Science Communication: How Motivated Reasoning and Identity Cues Amplify Opinion Polarization about Climate Mitigation Policies. *Communication Research*, 39(6):701–723.

Hazlett, C. and Mildenberger, M. (2020). Wildfire Exposure Increases Pro-Environment Voting within Democratic but Not Republican Areas. *American Political Science Review*, 114(4):1359–1365.

Healy, A. and Malhotra, N. (2009). Myopic Voters and Natural Disaster Policy. *American Political Science Review*, 103(3):387–406.

Healy, A. and Malhotra, N. (2010). Random Events, Economic Losses, and Retrospective Voting: Implications for Democratic Competence. *Quarterly Journal of Political Science*, 5(2):193–208.

Hoerling, M., Eischeid, J., Kumar, A., Leung, L. R., Mariotti, A., Mo, K., Schubert, S., and Seager, R. (2014). Causes and Predictability of the 2012 Great Plains Drought. *Bulletin of the American Meteorological Society*, 95(2):269–282.

Hooghe, M. and Dassonneville, R. (2018). Explaining the Trump Vote: The Effect of Racist Resentment and Anti-Immigrant Sentiments. *PS: Political Science & Politics*, 51(3):528–534.

Kahan, D. M., Peters, E., Wittlin, M., Slovic, P., Ouellette, L. L., Braman, D., and Mandel, G. (2012). The polarizing impact of science literacy and numeracy on perceived climate change risks. *Nature Climate Change*, 2(10):732–735.

Karakoc, D. B., Wang, J., and Konar, M. (2022). Food flows between counties in the United States from 2007 to 2017. *Environmental Research Letters*, 17(3):034035.

Key, V. O. (1966). *The Responsible Electorate: Rationality in Presidential Voting, 1936–1960*. Harvard University Press, Cambridge, MA.

Kramer, G. H. (1971). Short-Term Fluctuations in U.S. Voting Behavior, 1896–1964. *American Political Science Review*, 65(1):131–143.

Kuwayama, Y., Thompson, A., Bernknopf, R., Zaitchik, B., and Vail, P. (2019). Estimating the Impact of Drought on Agriculture Using the U.S. Drought Monitor. *American Journal of Agricultural Economics*, 101(1):193–210.

Lin, X., Ruess, P. J., Marston, L., and Konar, M. (2019). Food flows between counties in the United States. *Environmental Research Letters*, 14(8):084011.

McKee, S. C. (2008). Rural Voters and the Polarization of American Presidential Elections. *PS: Political Science & Politics*, 41(1):101–108.

Merkley, E. and Stecula, D. A. (2018). Party Elites or Manufactured Doubt? The Informational Context of Climate Change Polarization. *Science Communication*, 40(2):258–274.

Merkley, E. and Stecula, D. A. (2021). Party Cues in the News: Democratic Elites, Republican Backlash, and the Dynamics of Climate Skepticism. *British Journal of Political Science*, 51(4):1439–1456.

Mutz, D. C. (2018). Status threat, not economic hardship, explains the 2016 presidential vote. *Proceedings of the National Academy of Sciences*, 115(19):E4330–E4339.

Otkin, J. A., Anderson, M. C., Hain, C., Svoboda, M., Johnson, D., Mueller, R., Tadesse, T., Wardlow, B. D., and Brown, J. F. (2016). Assessing the evolution of soil moisture and vegetation conditions during the 2012 United States flash drought. *Agricultural and Forest Meteorology*, 218–219:230–242.

Palmer, W. C. (1965). Meteorological Drought. Research Paper 45, U.S. Weather Bureau, Washington, DC.

Scala, D. J. and Johnson, K. M. (2017). Political Polarization along the Rural–Urban Continuum? The Geography of the Presidential Vote, 2000–2016. *The ANNALS of the American Academy of Political and Social Science*, 672(1):162–184.

Schaffner, B. F., MacWilliams, M., and Nteta, T. (2018). Understanding White Polarization in the 2016 Vote for President: The Sobering Role of Racism and Sexism. *Political Science Quarterly*, 133(1):9–34.

Tesler, M. (2018). Elite Domination of Public Doubts about Climate Change (Not Evolution). *Political Communication*, 35(2):306–326.

Appendix

A. Theoretical framework

A.1 Belief updating

Let $f_M(D)$ and $f_O(D)$ denote the densities of recognized drought intensity $D \geq 0$ under the policy-relevant climate state M and the natural state O, respectively. Define the likelihood ratio

$$\ell(D) = \frac{f_M(D)}{f_O(D)}. \tag{A.1}$$

Assumption A.1 (Monotone likelihood ratio). The likelihood ratio is positive, continuously differentiable, and increasing in drought intensity:

$$\ell(D) > 0, \qquad \ell'(D) > 0. \tag{A.2}$$

Thus, more severe droughts provide stronger evidence for state M relative to state O. Let $\omega_0 > 0$ denote the common reference odds of state M relative to state O before the drought signal and the partisan adjustment are taken into account.

For a voter with partisan identity $\sigma \in \{D, R\}$, posterior beliefs satisfy

$$\ln\left(\frac{\mu_\sigma(D)}{1-\mu_\sigma(D)}\right) = \ln \omega_0 + \ln \ell(D) + m_\sigma(D). \tag{A.3}$$

Equivalently,

$$\mu_\sigma(D) = \frac{\omega_0\, \ell(D)\, e^{m_\sigma(D)}}{1 + \omega_0\, \ell(D)\, e^{m_\sigma(D)}}. \tag{A.4}$$

The partisan adjustments are

$$m_D(D) = \alpha_D + \gamma_D D, \qquad m_R(D) = -\alpha_R - \gamma_R D, \tag{A.5}$$

where $\alpha_D, \alpha_R, \gamma_D, \gamma_R \geq 0$. The level parameters α_D and α_R shift the assessment of state M at every drought intensity. The slope parameters γ_D and γ_R determine how the partisan adjustment changes as drought intensity increases. These adjustments are reduced-form belief wedges: they may represent differences in source trust, causal interpretation, policy relevance, or motivated processing. They do not require voters to disagree about whether the drought occurred.

A.2 From beliefs to platform evaluation

Let $V_P(S)$ denote the value, in state $S \in \{M, O\}$, of the platform offered by party $P \in \{D, R\}$. If a voter assigns probability q to state M, the expected relative value of the Democratic platform is

$$\Delta V(q) = q\,[V_D(M) - V_R(M)] + (1-q)\,[V_D(O) - V_R(O)]\,. \tag{A.6}$$

Define

$$\kappa \equiv [V_D(M) - V_R(M)] - [V_D(O) - V_R(O)] > 0. \tag{A.7}$$

The condition $\kappa > 0$ means that a stronger belief in state M makes the Democratic platform relatively more attractive. It does not require the Democratic platform to be preferred overall.

The change in the relative appeal of the Democratic platform caused by observing drought intensity D is

$$A_\sigma(D) \equiv \Delta V(\mu_\sigma(D)) - \Delta V(\mu_\sigma(0)) = \kappa\,[\mu_\sigma(D) - \mu_\sigma(0)]\,. \tag{A.8}$$

Lemma A.1 (Marginal response to drought intensity). Define the common informational content of a marginal increase in drought intensity as

$$s(D) \equiv \frac{d \ln \ell(D)}{dD} = \frac{\ell'(D)}{\ell(D)} > 0. \tag{A.9}$$

Then

$$\frac{dA_\sigma(D)}{dD} = \kappa\,\mu_\sigma(D)\,[1-\mu_\sigma(D)]\,[s(D) + m'_\sigma(D)]\,. \tag{A.10}$$

Consequently,

$$\operatorname{sign}\left(\frac{dA_\sigma(D)}{dD}\right) = \operatorname{sign}[s(D) + m'_\sigma(D)]\,. \tag{A.11}$$

Proof. Differentiating the log-odds equation gives

$$\frac{\mu'_\sigma(D)}{\mu_\sigma(D)\,[1-\mu_\sigma(D)]} = s(D) + m'_\sigma(D). \tag{A.12}$$

Therefore,

$$\mu'_\sigma(D) = \mu_\sigma(D)\,[1-\mu_\sigma(D)]\,[s(D) + m'_\sigma(D)]\,. \tag{A.13}$$

Differentiating the definition of $A_\sigma(D)$ and using $\kappa > 0$ establishes the result. The probability-scaling term $\mu_\sigma(D)\,[1-\mu_\sigma(D)]$ is strictly positive whenever posterior beliefs are interior. □

A.3 Partisan response patterns

For Democratic voters, $m'_D(D) = \gamma_D$, and therefore

$$\frac{dA_D(D)}{dD} = \kappa\,\mu_D(D)\,[1-\mu_D(D)]\,[s(D)+\gamma_D] > 0. \tag{A.14}$$

For Republican voters, $m'_R(D) = -\gamma_R$, and therefore

$$\frac{dA_R(D)}{dD} = \kappa\,\mu_R(D)\,[1-\mu_R(D)]\,[s(D)-\gamma_R]\,. \tag{A.15}$$

At a reference drought intensity $\overline{D}$, define the group-specific marginal sensitivities

$$\begin{aligned}
\delta_D &\equiv \left.\frac{dA_D(D)}{dD}\right|_{D=\overline{D}} = \kappa\,\mu_D(\overline{D})\left[1-\mu_D(\overline{D})\right]\left[s(\overline{D})+\gamma_D\right] > 0,\\
\delta_R &\equiv -\left.\frac{dA_R(D)}{dD}\right|_{D=\overline{D}} = \kappa\,\mu_R(\overline{D})\left[1-\mu_R(\overline{D})\right]\left[\gamma_R - s(\overline{D})\right].
\end{aligned} \tag{A.16}$$

The sign of δ_R is not imposed. The definition uses a minus sign so that $\delta_R > 0$ corresponds to a movement *away* from the Democratic platform among Republican voters.

Proposition A.1 (Partisan responses). At $D = \overline{D}$:

(i) Democratic voters necessarily move in the Democratic direction, because $\delta_D > 0$.

(ii) If $\gamma_R < s(\overline{D})$, then $\delta_R < 0$ and Republican voters also move in the Democratic direction; the common informational effect dominates the partisan adjustment.

(iii) If $\gamma_R = s(\overline{D})$, then $\delta_R = 0$ and the two forces exactly offset one another.

(iv) If $\gamma_R > s(\overline{D})$, then $\delta_R > 0$ and Republican voters move away from the Democratic platform; this is partisan backlash.

Proof. The Democratic result follows directly from $s(\overline{D}) > 0$, $\gamma_D \geq 0$, and the positive scaling terms in the first line of (5). The Republican result follows from the sign of $\gamma_R - s(\overline{D})$ in the second line.

The level parameters α_D and α_R do not enter these sign conditions directly. They can nevertheless affect response magnitudes through the posterior-variance term $\mu_\sigma(\overline{D})\left[1-\mu_\sigma(\overline{D})\right]$, which is largest near one-half and approaches zero as beliefs become degenerate.

A.4 Aggregation to counties

Let $\pi_i \in [0,1]$ denote county i's historical Democratic vote share, used as a proxy for the proportion of Democratic voters. A voter v with partisan identity σ votes Democratic with probability

$$P_{v\sigma}(D) = \Lambda[u_v + A_\sigma(D)]\,, \qquad \Lambda'(\cdot) > 0, \tag{A.17}$$

where u_v captures the voter's baseline relative preference and $\Lambda(\cdot)$ is an increasing voting-response function. County-level Democratic vote share is

$$y_i(D) = \pi_i \, \mathbb{E}_D[\Lambda(u_v + A_D(D))] + (1 - \pi_i) \, \mathbb{E}_R[\Lambda(u_v + A_R(D))] \, . \tag{A.18}$$

To keep the aggregation exact at the reference point, define voting-probability-adjusted sensitivities

$$\widetilde{\delta}_D \equiv \mathbb{E}_D\Big[\Lambda'\Big(u_v + A_D(\overline{D})\Big)\Big] \, \delta_D, \qquad \widetilde{\delta}_R \equiv \mathbb{E}_R\Big[\Lambda'\Big(u_v + A_R(\overline{D})\Big)\Big] \, \delta_R. \tag{A.19}$$

Because the expectations multiplying δ_D and δ_R are positive, this adjustment preserves their signs. To avoid notational clutter, suppress the tildes below and interpret δ_D and δ_R as the voting-probability-adjusted sensitivities.

Proposition A.2 (County-level marginal effect). The marginal effect of local drought on county i's Democratic vote share at the reference intensity is

$$\left.\frac{\partial y_i(D)}{\partial D}\right|_{D=\overline{D}} = \pi_i \delta_D - (1 - \pi_i) \, \delta_R = -\delta_R + (\delta_D + \delta_R) \, \pi_i \equiv \beta_0 + \beta_1 \pi_i, \tag{A.20}$$

where

$$\beta_0 = -\delta_R, \qquad \beta_1 = \delta_D + \delta_R. \tag{A.21}$$

Proof. Differentiate the county vote-share equation (A.18). The Democratic group contributes $\pi_i \delta_D$. Since $dA_R/dD = -\delta_R$ at the reference point, the Republican group contributes $-(1 - \pi_i) \, \delta_R$. Collecting terms yields the stated expression. □

The positive partisan gradient tested in the main text corresponds to the ordering condition

$$\beta_1 > 0 \quad \Longleftrightarrow \quad \delta_D > -\delta_R. \tag{A.22}$$

This condition is automatic under strict Republican backlash ($\delta_R > 0$). When Republican voters also move toward the Democratic platform ($\delta_R < 0$), it requires their pro-Democratic response to be weaker than the Democratic response. Thus, the primitive sign restrictions alone establish an unambiguously positive Democratic response and an ambiguous Republican response; the positive interaction coefficient is the additional empirical ordering implied by partisan attenuation.

Three benchmark cases follow.

Uniform response. If both groups move equally in the Democratic direction, then $-\delta_R = \delta_D > 0$. Hence $\beta_0 = \delta_D > 0$ and $\beta_1 = 0$: drought has the same positive effect in every county.

Attenuation without backlash. If $0 \leq -\delta_R < \delta_D$, Republican voters also move in the Democratic direction but less strongly than Democratic voters. Hence $\beta_0 \geq 0$ and $\beta_1 > 0$: the effect is positive but weaker in more Republican counties.

Strict backlash. If $\delta_R > 0$, then $\beta_0 < 0$ and $\beta_1 > 0$. The effect is negative in sufficiently Republican counties and crosses zero at

$$\pi_D^* = -\frac{\beta_0}{\beta_1} = \frac{\delta_R}{\delta_D + \delta_R} \in (0, 1). \tag{A.23}$$

For $\pi_i < \pi_D^*$, Republican backlash dominates and drought lowers Democratic vote share. For $\pi_i > \pi_D^*$, the positive response among Democratic voters dominates and drought raises Democratic vote share.

A.5 Network transmission

For flow type m, define county i's network exposure to drought as

$$N_{it}^m = \sum_{j \neq i} w_{ij}^m D_{jt}, \qquad w_{ij}^m \geq 0, \qquad \sum_{j \neq i} w_{ij}^m = 1. \tag{A.24}$$

The weights therefore make N_{it}^m a weighted average of drought intensity among county i's economic partners. Let the climate-relevant signal perceived in county i be

$$D_{it}^{\star,m} = D_{it} + \varphi N_{it}^m, \qquad \varphi \geq 0. \tag{A.25}$$

The parameter φ captures the political visibility or climate relevance of network-transmitted drought relative to local drought. The common-salience specification assumes that the same φ applies to Democratic and Republican voters and that local and network exposure enter the same attribution mechanism through $D_{it}^{\star,m}$.

Proposition A.3 (Common network salience). Under the common-salience specification, the marginal local and network effects at the reference signal satisfy

$$\frac{\partial y_{it}}{\partial D_{it}} = \beta_0 + \beta_1 \pi_i, \qquad \frac{\partial y_{it}}{\partial N_{it}^m} = \varphi \left(\beta_0 + \beta_1 \pi_i\right). \tag{A.26}$$

Equivalently, if the local coefficients are (β_{0D}, β_{1D}) and the network coefficients are $(\beta_{0N}^m, \beta_{1N}^m)$, then

$$\beta_{0N}^m = \varphi \beta_{0D}, \qquad \beta_{1N}^m = \varphi \beta_{1D}. \tag{A.27}$$

When $\varphi > 0$ and $\beta_{1D} \neq 0$, local and network effects change sign at the same partisan threshold:

$$\pi_{N^m}^\star = -\frac{\beta_{0N}^m}{\beta_{1N}^m} = -\frac{\beta_{0D}}{\beta_{1D}} = \pi_D^*. \tag{A.28}$$

If $\varphi = 0$, the network effect is identically zero and a network threshold is not defined.

Proof. Since $D_{it}^{\star,m} = D_{it} + \varphi N_{it}^{m}$, the chain rule gives $\partial D_{it}^{\star,m}/\partial D_{it} = 1$ and $\partial D_{it}^{\star,m}/\partial N_{it}^{m} = \varphi$. Applying Proposition A.2 to the perceived signal yields the two derivatives. Coefficient proportionality and threshold equality follow immediately when the relevant denominators are nonzero. □

A.5.1 Partisan-specific network salience

If network exposure is perceived with group-specific salience parameters φ_D and φ_R, the network marginal effect becomes

$$\frac{\partial y_{it}}{\partial N_{it}^{m}} = \pi_i \varphi_D \delta_D - (1 - \pi_i)\, \varphi_R \delta_R = -\varphi_R \delta_R + (\varphi_D \delta_D + \varphi_R \delta_R)\, \pi_i. \tag{A.29}$$

Under strict backlash and positive salience, its zero-crossing is

$$\pi_{N^m}^{\star} = \frac{\varphi_R \delta_R}{\varphi_D \delta_D + \varphi_R \delta_R}. \tag{A.30}$$

Hence local and network thresholds need not coincide. A threshold difference rejects equal salience across partisan groups, but not necessarily the broader attribution mechanism.

A.5.2 Direction of transmission

Let N_{it}^{S} denote exposure through supplier counties and N_{it}^{B} exposure through buyer counties. Define total and directional exposure as

$$N_{it} = N_{it}^{S} + N_{it}^{B}, \qquad N_{it}^{\Delta} = N_{it}^{S} - N_{it}^{B}. \tag{A.31}$$

If voters respond only to the total amount of network-transmitted drought, the perceived signal depends on N_{it} but not on N_{it}^{Δ}. Conditional on total exposure, the model therefore implies the exclusion restriction

$$\frac{\partial y_{it}}{\partial N_{it}^{\Delta}} = 0. \tag{A.32}$$

A nonzero directional effect would indicate an additional channel—for example, input-cost and production disruptions from supplier drought versus demand losses from buyer drought—rather than, by itself, disproving the broader climate-attribution mechanism.

A.6 First-order empirical representation and testable restrictions

A first-order approximation around the reference perceived drought signal, with constants absorbed by county and election fixed effects, gives the restricted empirical representation

$$y_{it} = \alpha_i + \lambda_t + (\beta_0 + \beta_1 \pi_i)\, D_{it} + \varphi\, (\beta_0 + \beta_1 \pi_i)\, N_{it}^m + u_{it}. \tag{A.33}$$

Here y_{it} is Democratic vote share, α_i denotes county fixed effects, λ_t denotes election fixed effects, and u_{it} collects unmodeled determinants of voting. The fixed effects absorb the constants generated by Taylor expansion and persistent or election-wide components of baseline political support.

The unrestricted specification used for testing can be written compactly as

$$y_{it} = \alpha_i + \lambda_t + L_{it} + G_{it}^m + H_{it}^{\Delta} + u_{it}, \tag{A.34}$$

where the three exposure components are

$$L_{it} = (\beta_{0D} + \beta_{1D}\pi_i)\, D_{it}, \qquad G_{it}^m = (\beta_{0N}^m + \beta_{1N}^m \pi_i)\, N_{it}^m, \qquad H_{it}^{\Delta} = \left(\beta_{0N}^{\Delta} + \beta_{1N}^{\Delta}\pi_i\right) N_{it}^{\Delta}. \tag{A.35}$$

The theory yields the following restrictions.

Partisan gradient. Under the partisan-ordering condition, the effect of both local and network drought becomes more favorable to Democrats as historical Democratic alignment increases:

$$\beta_{1D} > 0, \qquad \beta_{1N}^m > 0. \tag{A.36}$$

Under strict backlash, the effects at $\pi_i = 0$ are also negative:

$$\beta_{0D} < 0, \qquad \beta_{0N}^m < 0. \tag{A.37}$$

Common proportionality and equal thresholds. Common network salience implies

$$\frac{\beta_{0N}^m}{\beta_{0D}} = \frac{\beta_{1N}^m}{\beta_{1D}} = \varphi, \tag{A.38}$$

whenever the denominators are nonzero. A denominator-robust equivalent of equal proportionality—and therefore of equal zero-crossing thresholds—is

$$\beta_{0N}^m \beta_{1D} - \beta_{1N}^m \beta_{0D} = 0. \tag{A.39}$$

No separate directional effect. Once total network exposure is controlled for,

$$\beta_{0N}^{\Delta} = 0, \qquad \beta_{1N}^{\Delta} = 0. \tag{A.40}$$

The first restriction tests the central partisan-attribution gradient. The second tests whether local and network drought activate the same mechanism with a common attenuation factor. The third tests whether the direction of economic transmission creates an additional electoral channel.

B. Data Construction and Descriptive Evidence

B.1 Food-flow network construction and characteristics

Table B.1: Commodity composition of the 2007 food-flow network

Commodity group	Description	Share of selected flows (%)
SCTG 02	Cereal grains, including seed	71.2
SCTG 03	Other agricultural products, excluding animal feed, cereal grains, and forage products	17.0
SCTG 04	Animal feed, eggs, honey, and other products of animal origin	11.8
Total	Selected SCTG 02–04 flows	100.0

Notes: Shares are calculated from the mass of estimated county-to-county food flows in 2007. The network includes SCTG 02–04 and excludes within-county flows. Source: Karakoc et al. (2022).

B.2 Historical partisan alignment

Table B.2: Descriptive statistics for historical partisan alignment, by metro status

	Metropolitan	**Nonmetropolitan**	**Full sample**	**Difference**
Counts				
Number of counties	1,057	2,012	3,069	—
Central tendency and dispersion of π_i				
Mean	0.457	0.444	0.448	0.014
Median	0.454	0.439	0.443	0.015
Standard deviation	0.102	0.110	0.107	—
Overlap of distributions (quantiles of π_i)				
P10 → P90	0.335–0.586	0.306–0.590	0.317–0.589	—
P25 → P75	0.383–0.524	0.369–0.511	0.375–0.516	—
Min → Max	0.190–0.806	0.120–0.835	0.120–0.835	—
Democratic-aligned counties ($\pi_i > 0.5$)				
Share of counties	32.2%	28.4%	29.7%	3.8 pp
Metropolitan/nonmetropolitan separation				
Standardized difference of means	—	—	—	0.13

Notes: Metropolitan status is based on the 2003 Rural–Urban Continuum Codes. The "Difference" column reports metropolitan minus nonmetropolitan values; "pp" denotes percentage points. Differences use unrounded values.

B.3 Sample descriptives

Table B.3: Descriptive statistics for baseline regression variables

Variable	N	Mean	SD	Min	P25	Median	P75	Max
Local drought, D_{it}	15 345	0.717	1.000	0.000	0.000	0.077	1.260	6.286
Mean network exposure, N_{it}^{m}	15 345	0.698	0.838	0.000	0.050	0.347	1.087	5.157
Directional network exposure, N_{it}^{Δ}	15 345	0.009	0.731	−4.312	−0.162	0.000	0.161	4.505
Regional drought, D_{it}^{reg}	15 345	0.706	0.601	0.015	0.294	0.476	0.955	2.824

Note: The unit of observation is the county-election. D_{it} is standardized local drought; N_{it}^{m} and N_{it}^{Δ} are mean and directional food-network drought exposure, respectively; and D_{it}^{reg} is average drought among other counties within 1,000 km. Network measures use predetermined 2007 SCTG 02–04 flows.

C. Robustness

Table C.1: Marginal effects after excluding the 2012 and 2016 elections

Sample	Exposure channel	$\pi_i = 0.25$	$\pi_i = 0.50$	$\pi_i = 0.75$
Excluding 2012	Local drought, D	-0.387^{**}	0.049	0.485
		(0.183)	(0.202)	(0.445)
	Mean network exposure, N^m	-0.279	0.160	0.598
		(0.223)	(0.193)	(0.448)
	Directional network exposure, N^Δ	0.215	0.098	-0.019
		(0.174)	(0.110)	(0.293)
	Combined effect, $D + N^m$	-0.666^{**}	0.209	1.083^{*}
		(0.281)	(0.307)	(0.618)
Excluding 2016	Local drought, D	-0.297^{*}	0.048	0.393
		(0.169)	(0.171)	(0.361)
	Mean network exposure, N^m	-0.569^{***}	0.208	0.984^{***}
		(0.212)	(0.162)	(0.328)
	Directional network exposure, N^Δ	-0.132	-0.220^{*}	-0.308
		(0.156)	(0.130)	(0.250)
	Combined effect, $D + N^m$	-0.866^{***}	0.255	1.376^{***}
		(0.263)	(0.264)	(0.432)

*Note:*Marginal effects are from the all-channels specification and are evaluated at the indicated values of π_i. The combined effect corresponds to simultaneous one-unit increases in D and N^m. Exposure variables retain the full-sample normalization. Conley standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

C.1 Election-specific heterogeneity in 2012

To distinguish the loss of precision caused by excluding 2012 from genuine 2012-specific coefficient heterogeneity, we estimate a pooled specification that retains all five presidential elections. Let $J_t^{2012} = \mathbf{1}\{t = 2012\}$. We estimate

$$y_{it} = \alpha_i + \lambda_t + \varphi \pi_i J_t^{2012} + \sum_k X_{it}^k \left[\beta_{0k} + \beta_{1k}\pi_i + J_t^{2012}\left(\delta_{0k} + \delta_{1k}\pi_i\right)\right] + \theta D_{it}^{reg} + \eta \pi_i (t - \bar{t}) + \varepsilon_{it}, \tag{C.1}$$

where $k \in \{D, N^m, N^\Delta\}$. The term $\pi_i J_t^{2012}$ allows the relationship between historical Democratic support and the Democratic vote share to shift specifically in 2012, independently of drought exposure.

For exposure channel k, the marginal relationship is

$$\frac{\partial y_{it}}{\partial X_{it}^k} = \begin{cases} \beta_{0k} + \beta_{1k}\pi_i, & t \neq 2012, \\ (\beta_{0k} + \delta_{0k}) + (\beta_{1k} + \delta_{1k})\pi_i, & t = 2012. \end{cases} \tag{C.2}$$

The partisan gradient is therefore β_{1k} outside 2012 and $\beta_{1k} + \delta_{1k}$ in 2012. Equality of the partisan gradients is tested using

$$H_0^{\text{gradient}} : \delta_{1k} = 0 \tag{C.3}$$

A stronger test asks whether the entire exposure-response function is unchanged in 2012:

$$H_0^{\text{slope}} : \delta_{0k} = \delta_{1k} = 0 \tag{C.4}$$

Table C.2: Election-specific heterogeneity in 2012

	Outside 2012	2012	Difference	$W(q)$	p-value
Panel A. Estimated partisan gradients					
Local drought, D	1.907*	2.367**	0.459	–	0.759
	(1.033)	(1.116)	(1.499)		
Mean network exposure, N^m	2.054*	2.021	−0.033	–	0.987
	(1.113)	(1.428)	(1.951)		
Directional network exposure, N^Δ	0.054	−0.675	−0.729	–	0.381
	(0.813)	(0.638)	(0.832)		
Panel B. Wald tests of 2012-specific differences					
Local slope function	–	–	–	3.037 (2)	0.219
Mean-network slope function	–	–	–	0.677 (2)	0.713
Directional-network slope function	–	–	–	29.345 (2)	< 0.001
Local and mean-network partisan gradients	–	–	–	0.107 (2)	0.948
All partisan gradients	–	–	–	0.853 (3)	0.837
Local and mean-network slope functions	–	–	–	3.612 (4)	0.461
All slope functions	–	–	–	33.870 (6)	< 0.001
Observations			15,345		
Counties			3,069		
Elections			5		

Note: Panel A reports partisan gradients outside 2012, in 2012, and their differences. Panel B reports Wald tests of 2012-specific differences; a slope-function test jointly evaluates differences in the level coefficient and its interaction with π_i. Conley standard errors are reported in parentheses. W is the Wald statistic and q the number of restrictions. $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table C.3: Election-specific heterogeneity in 2012: full estimates

Variable	Coefficient	Standard error	p-value
Local drought, D_{it}	-0.907^{**}	0.403	0.024
Partisan gradient, $D_{it} \times \pi_i$	1.907^{*}	1.033	0.065
2012 level difference, $D_{it} \times J_t^{2012}$	0.115	0.571	0.840
2012 gradient difference, $D_{it} \times \pi_i \times J_t^{2012}$	0.459	1.499	0.759
Mean network exposure, N_{it}^{m}	-0.747	0.461	0.105
Partisan gradient, $N_{it}^{m} \times \pi_i$	2.054^{*}	1.113	0.065
2012 level difference, $N_{it}^{m} \times J_t^{2012}$	-0.202	0.768	0.792
2012 gradient difference, $N_{it}^{m} \times \pi_i \times J_t^{2012}$	-0.033	1.951	0.987
Directional network exposure, N_{it}^{Δ}	0.147	0.354	0.678
Partisan gradient, $N_{it}^{\Delta} \times \pi_i$	0.054	0.813	0.948
2012 level difference, $N_{it}^{\Delta} \times J_t^{2012}$	-0.332	0.358	0.353
2012 gradient difference, $N_{it}^{\Delta} \times \pi_i \times J_t^{2012}$	-0.729	0.832	0.381
Regional drought, D_{it}^{reg}	0.531	0.513	0.300
Partisan trend, $\pi_i \times (t - \bar{t})$	-1.114^{***}	0.225	< 0.001
2012-specific partisan shift, $\pi_i \times J_t^{2012}$	2.362	3.153	0.454
Wald tests of 2012-specific differences			
Test	W	q	p-value
Local slope function	3.037	2	0.219
Mean-network slope function	0.677	2	0.713
Directional-network slope function	29.345	2	< 0.001
Local and mean-network partisan gradients	0.107	2	0.948
All partisan gradients	0.853	3	0.837
Local and mean-network slope functions	3.612	4	0.461
All slope functions	33.870	6	< 0.001
Model statistics			
Observations	15,345		
Counties	3,069		
Elections	5		
R^2	0.953		
County fixed effects	Yes		
Election fixed effects	Yes		

Notes: The model uses all five presidential elections, with $J_t^{2012} = 1\{t = 2012\}$. Conley standard errors are reported in parentheses. The table also reports Wald tests of 2012-specific coefficient differences; a slope-function test jointly evaluates the level and partisan-gradient differences. W is the Wald statistic and q the number of restrictions. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.