

Free Trade Attitudes and the Vote

Todd Hertling^{1‡}
James Bisbee¹

¹Department of Political Science, Vanderbilt University

[‡]To whom correspondence should be addressed: todd.r.hertling@vanderbilt.edu

October 26, 2023

Abstract

Does exposure to free trade’s negative consequences cause Americans to be more protectionist, and do those trade attitudes in turn influence their voting behavior? The field of international political economy (IPE) has long posited that economic self-interest governs individuals’ policy preferences in a theoretical framework known as open economy politics (OEP). Yet we know from individual-level data from a range of other subfields that political attitudes are influenced by far more than simple self-interest. We devise an empirical confrontation between the rational actor assumption undergirding the OEP framework and an alternative set of theories of political behavior which we combine under the umbrella term of partisan motivated reasoning (PMR). This study contributes to the field by combining individual-level panel data from the Views of Electoral Research (VOTER) study with firm-level measures of trade-related layoffs, derived from applications to the federal Trade Adjustment Assistance (TAA) program. We causally link exposure to free trade’s negative consequences to two outcomes of interest to the political science literature: opinions on free trade policies, and vote choice for president in the 2012, 2016, and 2020 U.S. presidential elections. We demonstrate that—while protectionist attitudes increase in response to increasing exposure to trade-related layoffs—these changes have very little, if any, relationship to their voting behavior, even in the 2016 U.S. presidential election, which was characterized by an especially high national salience for the issue of trade policy.

Keywords: open economy politics, partisan motivated reasoning, trade attitudes and voting behavior

Introduction

Between 2010 and 2020, the electoral successes of populist leaders across advanced industrial democracies threatened what had become a taken-for-granted consensus on the durability and desirability of liberal globalist policies. The stark anti-globalist platforms on which these populists ran villainized an internationalist elite who had neglected the material welfare of its own people for too long. Furthermore, the plausibility of an anti-trade platform was bolstered by a flurry of important economics research in the early 2010s that suggested free trade’s negative consequences spilled over beyond the specific workers in import-competing firms or industries to threaten the welfare of their towns and communities. Economics research of the time presented a tidy and persuasive explanation for the regional patterns of the anti-globalist wave of the 2010s: rationally self-interested voters were simply behaving in accordance with their rational self-interest.

This simple explanation was grounded in empirical patterns that were consistent with a theoretical framework that had become the dominant paradigm of international political economy (IPE) scholarship over the preceding 30 years: the open economy politics (OEP) model. Evidence for the populist phenomenon was found in the geographic patterns of support for populist candidates, who effectively drew support from the hardest-hit regions of the United States (Autor et al., 2016), Italy (Barone and Kreuter, 2021), England (Inglehart and Norris, 2016), Germany (Dippel, Gold and Heblich, 2015), and across western Europe writ large (Colantone and Stanig, 2018*a,b*). The OEP framework, which helps inform how economic interests interact with political institutions, holds that individual interests—at least as far as they pertain to the high politics of foreign policy—are best described by a rational actor model in which an individual’s welfare is determined by her position in the (local) labor market. The political disturbances of the 2010s, centered as they were in economically depressed areas hardest hit by free trade, were the natural manifestation of neglected interests.

However, individual-level evidence linking economic outcomes to political attitudes was harder to find, and at a theoretical level faced a strong challenge from decades of public opinion research that questioned, if not outright rejected, the assumption of rational self interest that lies at the heart of the OEP paradigm. In fact, a rich body of research, primarily found in the American politics literature, posits that individual attitudes are poorly explained by the rational actor model of individual behavior on which the OEP framework rests. Instead of leaders adjusting foreign policy subject to the interplay between interests and institutions, this competing explanation argues that public policy is the *cause* of individual policy preferences, that constituents loyally follow their co-partisan leaders, and that the causal sequencing from welfare outcomes to policy attitudes is—at best—dormant, if not outright reversed. Furthermore, this literature argues that almost all of the variation in voting behavior is explained not by policy preferences but rather by party identity. In short, this body of research challenges the microfounding assumption of OEP at its most fundamental level.

In this paper, we devise an empirical confrontation between these two understandings of political behavior and document evidence of the erosion of OEP’s microfoundations over the past decade of U.S. politics. We show that the evolution of attitudes on free trade in the United States between 2011 and 2020 is better explained—and increasingly so—by an alternative explanation of political behavior, referred to as partisan motivated reasoning (PMR). Our analysis does not imply that the debate between these theoretical understandings is zero-sum. Rather, our results describe crucial scope conditions under which the OEP and PMR frameworks are relatively better suited than the other at explaining the evolution of trade attitudes and behaviors over the last ten years. Our research provides an important challenge to the dominant explanation for the populist backlash against globalism and trade in modern democracies.

1 Our challenge to a growing consensus

Economic self-interest drives American trade preferences, or so the story goes. The common treatment of open economy politics (OEP) posits that self-interest drives citizens’ policy preferences, and this would seemingly translate to individual or group affinity for protectionism or free trade (Scheve and Slaughter, 2001). This approach is largely predicated on the idea that free trade benefits all countries when certain groups, or sectors of the economy, have a comparative advantage across particular factors of production—generally, either capital or labor (Hainmueller and Hiscox, 2006). Under the Stolper-Samuelson paradigm, for example, high-skilled workers in advanced industrialized economies would benefit more from open markets than the low-skilled workers because the former would own the abundant factors of production for which their economy holds an advantage, whereas the latter would not (Mansfield and Mutz, 2009; Hainmueller and Hiscox, 2006). These producer identities determine material welfare outcomes under either liberal or protectionist trade policy regimes, which in turn govern the relevant policy preferences that each group adopts, as demonstrated by the negative correlation between education and protectionism (Scheve and Slaughter, 2001). Alternatively, competing economic models differentiate winners and losers along the cleavage of industry, generating different predictions about who these groups are, but agreeing with the microfounding assumption that they adopt policy preferences in accordance with their rationally self-interested calculations (Hiscox, 2001). More recent innovations have expanded the definition of winners and losers beyond the individual to recognize that simply living close to a shuttered factory might also consign one to the “losers” category due to economic agglomeration and negative spillovers (Autor et al., 2016; Malgouyres, 2017; Colantone and Stanig, 2018*a,b*; Rickard, 2020). Still, the presumption that material self-interest can (1) be linked with international trade policy and (2) explain variation in related policy preferences is the modern “hard core” of the OEP framework (Lake, 2009).

Increasingly, however, the apparent dominance of OEP has been questioned. A behav-

ioral revolution in IPE argued that policy preferences on trade are a function not of material self-interest but rather non-economic variables including nationalism (Margalit, 2012), identity (Scheve, Stasavage et al., 2006), and gender (Guisinger, 2016). Hainmueller and Hiscox (2006) contend that one’s trade preferences are less a function of calculated self-interest than they are reflective of higher-level exposure to economic ideas about globalization. More broadly, Mansfield and Mutz (2009) suggest that citizens have challenges connecting policy with their own welfare; therefore, self-interest lacks the ability to shape issue attitudes. Perhaps the strongest challenge comes from Rho and Tomz (2017), who contend that most citizens are ignorant of the economic consequences that protectionist policies forebear. The authors conclude that “nonmaterial concerns including nationalism, ethnocentrism, racism, and ideology could overshadow if not override economic costs and benefits,” and they are correct to hint at the influence of social identity on trade preference (Rho and Tomz, 2017, p. 91). Where they leave room for exploration, however, is in the role that partisanship might play as an identity driving trade policy preferences.

Since the seminal publication of *The American Voter* in 1960, it has been difficult to challenge the explanatory power of party identification, its deep psychological impact on citizens, and the social identities that derive from it (Campbell et al., 1960). Retrospective assessments of elected leaders—combined with personal experiences—offer potential insight into the identities formed around partisanship, and they are most impactful when considering performance of the economy or a president’s handling of a war (Fiorina, 1981; Lenz, 2013; Vavreck, 2009; Sides, Tausanovitch and Vavreck, 2022). While a qualifier on the power of party identification comes in the form of historical context, particularly when social and moral issues like race and religion are involved (Highton and Kam, 2011; Abramowitz and Webster, 2016; Achen and Bartels, 2017), it is increasingly difficult to divorce the policy issue from the predisposition for partisan identity. The predictive power of partisan identity reverses the causal logic undergirding the OEP framework: instead of trade preferences being determined by one’s labor market position, they are adjusted to avoid dissonance with one’s

party platform, a phenomenon known as partisan motivated reasoning (PMR).¹

If party loyalty does play an out-sized role with issue opinion formation, the grip that the partisan lens holds on trade policy preferences should be no different. Yet, there is not much conclusive evidence on the degree to which trade attitudes weigh on vote choice at the individual level. Certainly, the retrospective voting literature offers insight into the potential power that economically-oriented assessments have on the vote in the aggregate, but citizens do not always properly align their perceptions of economic performance with reality (Vavreck, 2009; Mansfield and Mutz, 2009). Furthermore, trade is a far more complex concept, with little evidence to suggest that individuals understand the underlying economic concepts (Rho and Tomz, 2017), making it more likely that they will rely on partisan cues to guide their attitudes (Bullock, 2020). More precisely linking the electorate’s exposure to import competition and job loss at the local level is important to understanding trade preferences and vote choice.

In aggregate data, empirical support for the OEP framework is abundant, particularly post-2000. For example, Margalit (2011) explores whether globalization’s effects on the labor market impact U.S. voting behavior, measured at the county level. His findings show that job losses from trade—operationalized as applications to the Department of Labor’s Trade Adjustment and Assistance program (TAA)—may result in negative retrospective presidential voting patterns. Similarly, Autor et al. (2020) apply a different measure of commuting zone-level import competition to demonstrate that economic hardship associated with trade predicts political polarization writ large, and former President Donald Trump’s

¹We use PMR as an umbrella term that accommodates various theoretical strands explaining how individuals process information. These include bounded rationality perspectives, where party cues serve as a useful heuristic on which voters rely, as well as social identity theory, which assumes that individuals are biased toward identity-preservation. While these are critical to a deeper understanding of the microfoundations of individual attitudes on trade, they are beyond the scope of this project, focused on testing the stronger claim of strictly rational behavior.

2016 election in particular.

In sum, one might assume that the OEP framework is on solid theoretical ground. Even if education captures metropolitanism and socialization instead of skill endowments, it still produces empirical patterns that are well-understood by the OEP framework. And even if individuals are not sophisticated enough to understand the precise mechanisms by which their labor market position benefits from, or is threatened by, international trade, in aggregate they behave as though they do. Yet this body of work has largely ignored the American politics elephant in the room: PMR. In the following section, we describe the OEP versus PMR debate in the context of existing theories of preference formation and voting behavior to highlight where these frameworks diverge, which motivates the empirical analyses that follows.

2 Theoretical intuition

The rational actor account of how free trade might influence both attitudes on trade policy and subsequent voting behavior is straightforward. Free trade produces both positive and negative outcomes with an uneven distribution across the citizenry. While the specific identities of free trade's winners and losers are a matter of persistent economic debate, it seems clear that these groups exist.

Standard economic theory tells us that the winners comprise both diffuse consumers who benefit from the lower prices resulting from international competition and producers in exporting industries whose goods command higher prices on the world market than they would otherwise at home (Ricardo, 1821), or those with skills that are relatively abundant at home (Stolper and Samuelson, 1941). The losers are workers in import-competing industries, or those endowed with skills that are relatively scarce at home. They face lower wages and the prospect of job loss as domestic firms either move production overseas or go out of

business entirely when faced with the cheaper goods coming from abroad.

To this standard economist’s account, we might also add those whose economic fates are linked to the aforementioned winners and losers by proximity. This proximity might be spatial, as exemplified by the booms and busts of local real estate confronted with the shuttering of import competing factories and the new development of exporting firms (Autor, Dorn and Hanson, 2013; Pierce and Schott, 2016), or the erosion of local public goods and services (Feler and Senses, 2017). But the concept of proximity might also be industrial, as upstream suppliers with no direct foreign competitors nevertheless find their margins thinning when downstream import-competing buyers are forced out of business. Further still, it might be occupational: workers who have neither geographic proximity nor share the same industry as import competition may still anticipate a bleak future wherein their skills are no longer valued by a labor market that effectively innovates to stay relevant in the dynamic global economy (Bisbee and Rosendorff, 2020; Deng et al., 2022).

Determining the precise contours of these groups has been the focus of economic and IPE research for most of their fields’ existences, and it has certainly governed the lion’s share of scholarly attention over the last fifty years. However, attention to who wins and loses ignores an essential political research question: do individuals adopt policy positions consistent with their self-interest in the first place? PMR predicts that individuals develop policy preferences not as a function of rational self-interest, but as a function of partisan identity and elite cues. Thus, attitudes on trade policy are not a function of one’s own experiences with positive and negative trade outcomes, but of whether one’s co-partisan elites support or oppose free trade.

Implicit in the preceding discussion is a causal intuition about the OEP and PMR frameworks. Consider the causal diagram displayed in Figure 1, where the top panel (OEP) captures the more traditional economic paradigm. Here, the material consequences of trade (“Facts”) determine the winners and losers, who subsequently update their attitudes about

the value of free trade (“Attitudes”). Individuals reference these preferences when evaluating the party platforms of competing political candidates, who describe their issue positions with “cues.” Then, they select a candidate (“Vote”). Taking OEP to the extreme, partisanship should not matter at all.

Conversely, the bottom panel (PMR) summarizes the partisan motivated reasoning theory in which an individual’s partisanship (“Party ID”) is adopted early in life through socialization and is a core part of their identity. This identity then predominates candidate choice (“Vote”), and the information about that candidate’s policy positions (“Cues”), which in turn informs the individual’s attitudes. Taking PMR to the extreme, facts should have no role.

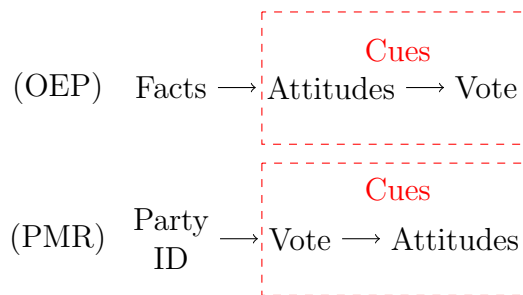


Figure 1: (OEP): A stylized depiction of the open economy politics framework. Voters update attitudes in response to facts, which in turn determine who they vote for on the basis of the cues parties give signaling their policy platforms. (PMR): A stylized depiction of partisan motivated reasoning in which a voter’s party ID determines their who they vote for, the cues from whom influence her policy preferences.

We propose adjudicating between these competing frameworks by focusing on two outcomes: policy preferences (“Attitudes”) and vote choice (“Vote”). If trade attitudes are best predicted by exposure to the material consequences of trade, and these attitudes predict vote choice, we would conclude that the OEP framework better explains the political outcomes of interest in the domain of trade policy. Conversely, if changes in individual trade attitudes are uncorrelated with changes in the material consequences of trade, and if instead these attitudes shift to match changing partisan positions on the topic, we would conclude that the PMR framework is more useful.

To guide our thinking on the determinants of these outcomes, we adapt the well-known Bayesian model of preference formation from the public opinion literature (Zechman, 1979; Achen, 1992; Bartels, 1993; Druckman and McGrath, 2019). This framework defines an individual’s beliefs in an optimal policy (i.e., trade) as a function of their priors on this issue, defined formally as $\pi_i(\mu) \sim \mathcal{N}(\hat{\mu}_{i,0}, \hat{\sigma}_{i,0}^2)$, and a new signal, defined as $x \sim \mathcal{N}(\mu_x, \hat{\sigma}_{i,x}^2)$. The parameter $\hat{\sigma}^2$ captures the variance associated with either the prior or the signal, and $\frac{1}{\hat{\sigma}^2}$ can be substantively interpreted as either the “certainty” of the individual’s prior, or the “precision” of the signal (Bartels, 1993) (or alternatively the “credibility” (Druckman and McGrath, 2019)). The individual’s posterior belief is thus:

$$\pi_i(\mu|x) \sim \mathcal{N}\left(\hat{\mu}_{i,0} + (x - \hat{\mu}_{i,0}) \left(\frac{\hat{\sigma}_{i,0}^2}{\hat{\sigma}_{i,0}^2 + \hat{\sigma}_{i,x}^2}\right), \frac{\hat{\sigma}_{i,0}^2 \hat{\sigma}_{i,x}^2}{\hat{\sigma}_{i,0}^2 + \hat{\sigma}_{i,x}^2}\right)$$

The flexibility of this framework is advantageous: new signals can be anything ranging from personal or proximate trade-related layoffs, to in- or out-party elite cues. As such, adjudicating between the OEP and PMR perspectives involves simply determining which signals are more prognostic of policy preferences. In the extreme, an OEP world is one in which $\sigma_{\text{facts}}^2 \ll \sigma_{\text{cues}}^2$ and $\sigma_{\text{Dem cues}}^2 = \sigma_{\text{Rep cues}}^2$, meaning that facts are more credible than elite cues, which are paid little heed, regardless of the partisanship of the cue giver. Conversely, a PMR world predicts that $\sigma_{\text{in-party cues}}^2 \ll \sigma_{\text{out-party cues}}^2 \leq \sigma_{\text{facts}}^2$, meaning that people only follow what in-party elites say, and either don’t perceive or don’t parse facts. Figure 2 visualizes these expectations, underscoring that—assuming the only signals include the negative consequences of free trade—we might nevertheless determine whether we are in the OEP or PMR setting by comparing the extent to which Democrats and Republicans respond to trade-related layoffs and elite cues from former President Donald Trump. In the OEP world, we would expect similar reactions regardless of partisanship, whereas in the PMR world, we would expect little to no response among Democrats, but a strong response among Republicans. (A third pattern would be that partisans diverge due to a backlash

effect among Democrats, possibly reflecting a world in which Trump’s protectionist cues provoke pro-free trade cues among Democrat elites.)

Bayesian belief formation

Differences in posterior due to differences in credibility

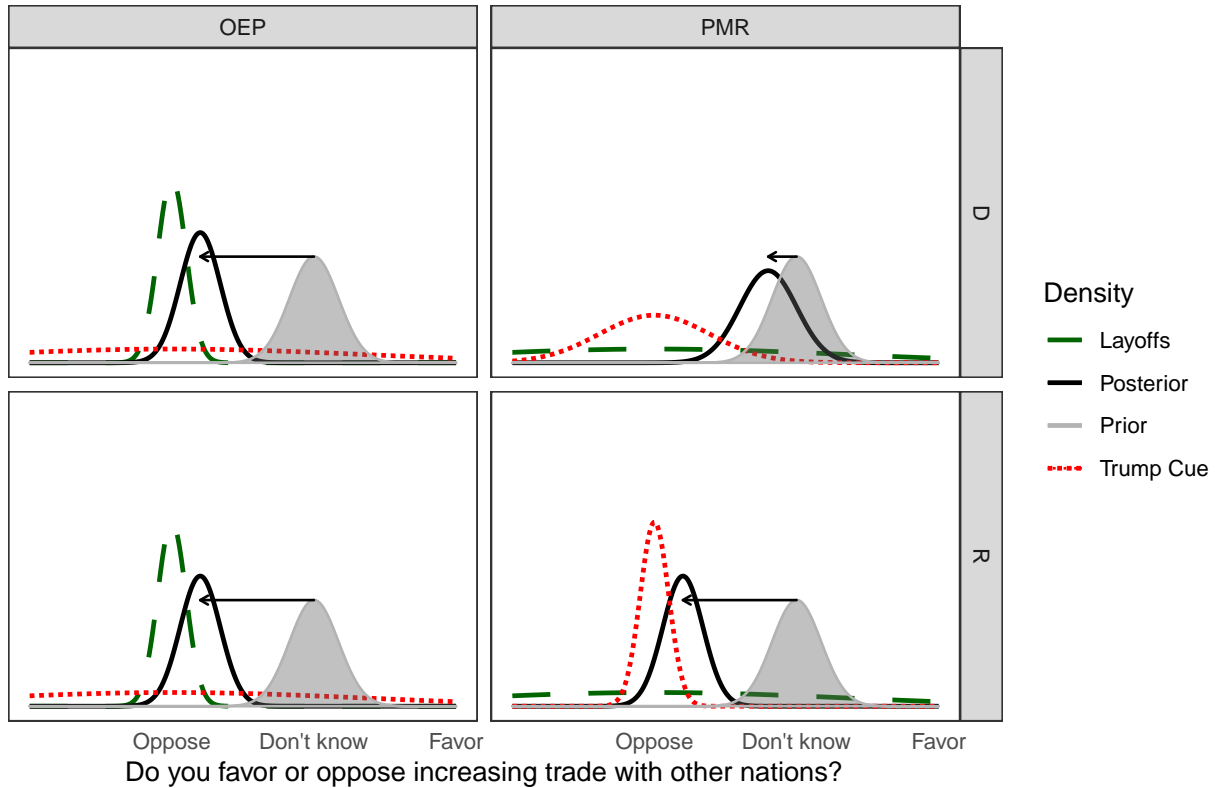


Figure 2: Example of partisan motivated reasoning in a Bayesian model of preference formation, with policy preferences indicated on the x-axis. Grey densities are an individual’s prior, which is the same across all facets. Solid black lines indicate the individual’s posterior, which is a function of either facts (“layoffs”, dashed green lines) or cues (“Trump cue”, dotted red lines). The left-column illustrates how preferences are theorized in the OEP setting, where elite cues are inconsequential while layoffs are equally influential for both Democrats (top row) and Republicans (bottom row). The right column illustrates how preferences are theorized in the PMR setting, where layoffs are inconsequential and Trump’s cues matter far more for Republicans than Democrats.

3 Data and Methods

Adjudicating between the OEP and PMR frameworks is complicated by empirical challenges, although the Bayesian framework helps provide some structure. Predicting vote choice with trade attitudes may be confounded by reverse causality if the PMR story is accurate. Predicting trade attitudes with material self-interest is hampered by either selection bias and the ecological inference fallacy (for evidence that relies on geographically aggregated data) or the construct validity of individual-level proxies for self-interest.² To overcome these challenges, we rely on data from the Views of Electoral Research (VOTER), a panel survey of 5,900 adults in the American electorate administered by the Democracy Fund Voter Study Group.³

We observe individuals' attitudes on trade at three points in time over this period: 2011, 2016, and 2020. We also have records of their (self-reported) vote choice in the 2008, 2012, 2016, and 2020 presidential elections. As discussed above, these two measures constitute our primary outcome or explanatory variable of interest. Each might be the outcome or the explanatory variable, depending on whether the OEP or PMR frameworks better describe the politics of trade over this period. We visualize the changes in both measures of interest in the below Figure 3, highlighting the higher variation in trade attitudes than in vote choice. Only 63%, 31%, and 22% of individuals supporting, not knowing, or opposing free trade retain those preference over the period of analysis, while these numbers rise to

²In the case of ecological inference fallacy, problematic is the drawing of causal inferences for the individual that derive from the aggregation of those individuals to geographical units like counties or congressional districts. Regarding construct validity, potential issues arise from assuming that education is an adequate proxy for skill or in attempting to predict attitudes on trade policy as a function of other questions about free trade. See Hainmueller and Hiscox 2006 for a discussion about the construct validity of education as a proxy for skill.

³Democracy Fund Voter Study Group. VIEWS OF THE ELECTORATE RESEARCH SURVEY, December 2021. Washington, D.C. <https://www.voterstudygroup.org/>.

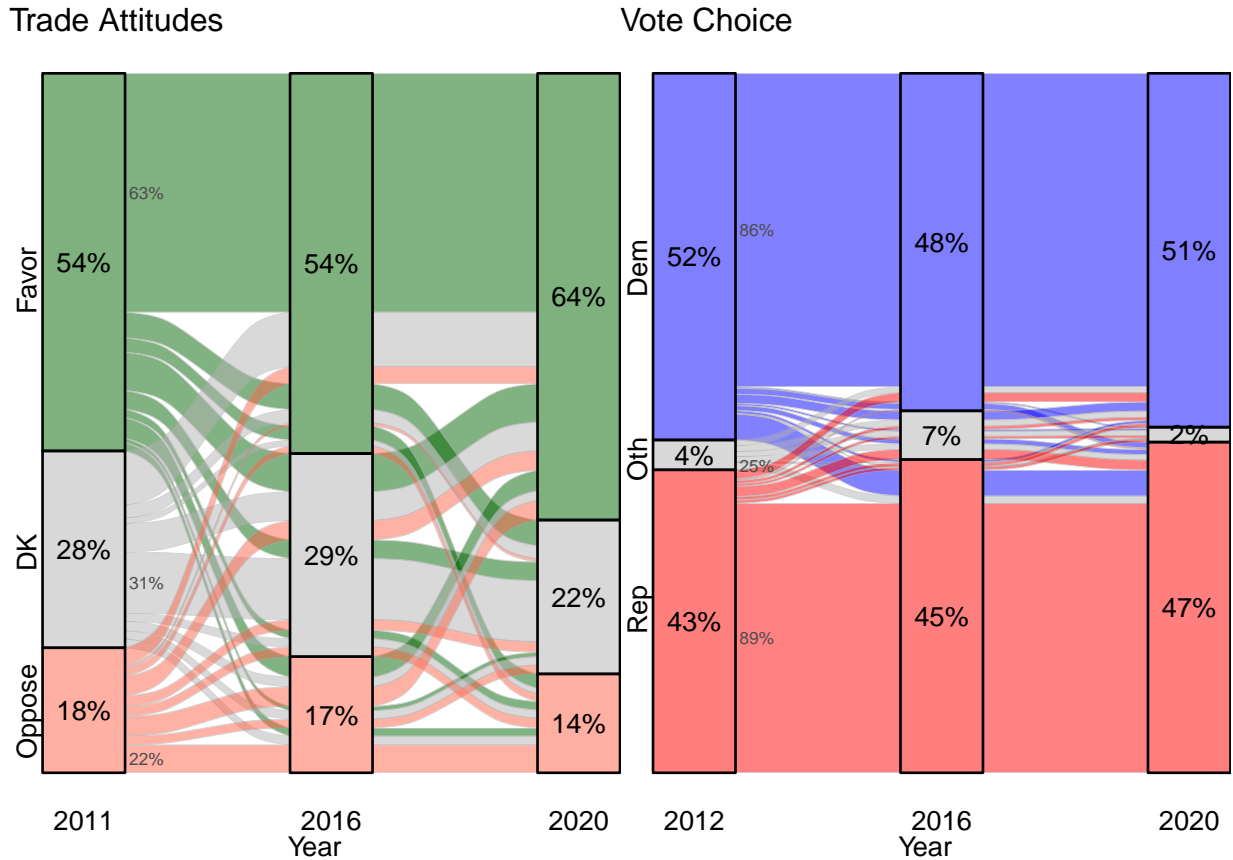


Figure 3: Flow diagrams displaying variation in main variables of interest. Left panel depicts changes in attitudes on trade in between 2011 and 2020. Right panel depicts changes in presidential vote choice between 2012 and 2020. Largest outflows in first period are indicated with light gray text.

86%, 25%, and 89% when it comes to stability in voting for Democrats, others, and Republicans. Additionally, we are interested in two other explanatory variables: free trade's negative consequences (the "facts" from Figure 1) and elite cues. For the former, we calculate the approximate latitude and longitude of each respondent based on his zip code, which allows us to characterize the degree of his exposure to trade-related layoffs, measured using applications to the Trade Adjustment Assistance (TAA) program (Margalit, 2011). Each application includes an estimated number of workers laid off as well as the specific location where they lost their job. To match survey respondents with these signals about free trade's negative consequences, we count the number of applications that occurred within an m -mile

radius over the preceding k days from when the respondent participated in the survey wave. An example of several different-sized radii, aggregated over periods of 90, 180, and 365 days is visualized for a single respondent in Figure 4.⁴

For the latter explanatory variable of elite cues, we exploit the fact that, over the course of this period, both the salience of and party positions on free trade changed dramatically. In 2011, former President Obama had yet to shift his attention to TTIP, and neither party had made trade policy a centerpiece of its platform for more than a decade. The parties were stereotyped on trade, as they had been since the early 1970s: Democrats were the party of labor and protectionism, while Republicans were the party of capital and free trade (Karol, 2009). Neither had made the issue a campaign focus for decades. By 2016, however, the landscape had shifted dramatically, with Republicans following then-candidate Donald Trump’s successful presidential campaign, which made protectionism a main talking point. We make the assumption that these nationwide signals were well-understood. Furthermore, we argue that the 2011 wave was fielded in a context of relatively weak partisan cues on trade, whereas the 2016 and 2020 waves were fielded in contexts of greater salience and clarity of party positions.

3.1 Methods

The panel nature of our data, measured at the individual level, enables us to evaluate the relative explanatory power of the OEP and PMR frameworks. By analyzing data measured at the individual level, we avoid the ecological inference challenges associated with work that measures outcomes at the level of the county, congressional district, or commuting zone. By measuring the same individuals over time, we can untangle questions of causal direction that plague the OEP versus PMR debate. Additionally, by restricting attention to only those

⁴Please see the SI for a discussion of the construct validity of this measure as a proxy for signals about the negative consequences of free trade.

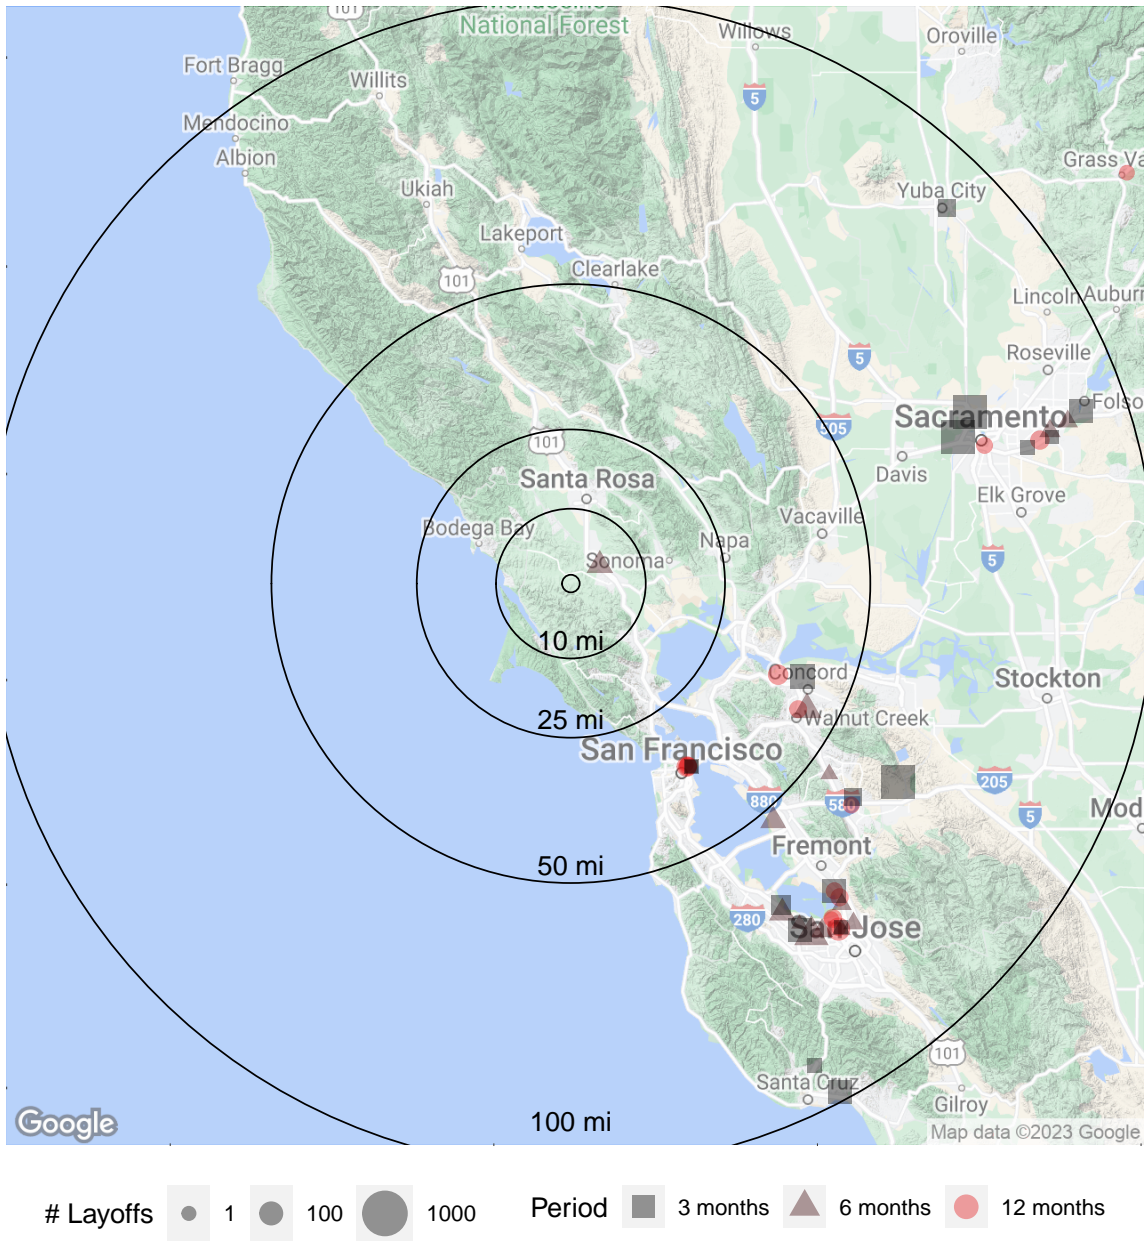


Figure 4: Example of exposure calculation for different radii for a respondent living north of San Francisco in 2011. X's indicate locations where trade-related layoffs occurred (measured in the TAA data), sized by the estimated number of affected workers and colored by whether the layoffs occurred 3 months prior to the survey (black), 6 months prior (dark red), or 12 months prior (red).

respondents who do not move zip codes between the survey waves, we reduce concerns about selection.

Seizing on the opportunity resident in this rich data, we rely on two related empirical methods for testing whether attitudes on trade are prognostic of vote choice. First, we implement an individual-fixed effects specification with the data organized by respondent-wave. Here, we predict the trade policy preferences or presidential vote choice y of respondent i in year t as a function of one’s exposure to trade-related layoffs TAA , along with lagged measures of the outcome. Crucially, the individual fixed effect α_i ensures the identification of variation in the outcome that occurs within an individual respondent, a measure of *change* in attitudes between waves as a function of the *change* in exposure. We also account for secular changes in the outcome via year fixed effects δ_t . Formally:

$$y_{it} = \alpha_i + \delta_t + \beta_1 \log\left(\sum_{t \in \{t-k, t\}} \sum_{d \in \{0, D\}} TAA_{itd}\right) + \varepsilon_{it} \quad (1)$$

In this specification, β_1 captures the effect of exposure to free trade’s negative consequences as proxied by the number of TAA applications within a certain distance of the respondent i over the preceding k years. Under the assumption that, conditional on the respondent’s prior outcome there are no remaining time varying confounders, the individual fixed effects allow us to interpret β_1 as a causally identified estimate of the effect of exposure on either attitudes or vote choice. To address the concern that these layoffs may be correlated with other place-based time-varying confounders, we run a placebo test in which we predict previous attitudes as a function of future layoffs.

Our second specification borrows from Lenz (2013) and organizes the data with each row indicating a respondent, and each column indexed by the year in which the measure was recorded. Here, we are only interested in the main measures of interest—attitudes and votes—and we predict each as a function of the other, controlling for the lagged measure of the dependent variable. Formally,

$$pref_{i,2016} = \alpha_g + \beta_1 pref_{i,2011} + \beta_2 vote_{i,2012} + \varepsilon_i \quad (2)$$

where α_g are district fixed effects, $pref_{i,t}$ denotes respondent i 's attitudes on trade in year t , and $vote_{i,t}$ denotes their vote choice. We are particularly interested in the β_2 coefficient, which captures the conditional relationship between respondent i 's 2012 vote choice and her 2016 attitudes on trade. We then compare to the analogous ρ_2 coefficient estimated in:

$$vote_{i,2016} = \alpha_g + \rho_1 vote_{i,2012} + \rho_2 pref_{i,2011} + \varepsilon_i \quad (3)$$

Substantively, $\beta_2 \gg \rho_2$ suggests that an individual's 2012 vote choice is far more prognostic of her 2016 attitude on trade than is her 2011 attitude on trade prognostic of her 2016 vote choice. Theoretically, such a pattern would be consistent with the PMR expectation that her vote determines her attitude, rather than the reverse, which would be consistent with the OEP expectation.

4 Results

4.1 Attitudes as a function of exposure

A first order question in the causal chain linking free trade's negative consequences to politics is whether individuals perceive these consequences in the first place. Existing research by Rho and Tomz (2017) argues that individuals are unable to connect even the most simple economic accounts of free trade's winners and losers to themselves. Conversely, there is abundant evidence linking geographically aggregated exposure to free trade's negative consequences with politically relevant outcomes (Margalit, 2011; Dippel, Gold and Heblich, 2015; Autor et al., 2016, 2017; Colantone and Stanig, 2018*a,b*).

To test, we start by predicting an individual's attitude on free trade as a function of their exposure to free trade's negative consequences using the specification described in

Equation 1 above. Our main results map these layoffs to each individual by summing up TAA applications that occurred within a 200-mile radius of the respondent over the 4 months prior to when one was surveyed, dropping applications that were rejected or terminated, filed on behalf of workers by unions or state officials, or were certified on the basis of upstream impacts of free trade.⁵ This measure is skewed, prompting us to log its value plus 1, although we also test the robustness of our conclusions to raw and binarized versions of the exposure measure. All regressions employ the panel weights provided by the VOTER data, and standard errors are clustered on the respondent. We drop any respondents who did not appear in all three waves, as well as those who moved within the period of the sample, yielding a final sample of 2,314.⁶

Each regression uses a linear probability model where the outcome is 1 if the respondent indicated that he favors expanding imports (column 1), 1 if the respondent indicated that he doesn't know (column 2), or 1 if the respondent indicated that he opposes expanding imports (column 3). As illustrated in Table 1, there is statistically and substantively significant support for the conclusion that exposure to trade-related layoffs is associated with a decline in favoring expanded imports by roughly three percentage points, against a mean of between 54% and 65% (see Figure 3 above). We also find a roughly 2.5 percentage point increase in opposition to expanding imports (column 2). There is no evidence of a systematic relationship between exposure and uncertainty over where one stands on the issue of trade (column 3).⁷

⁵We attempted to isolate the trade-related layoffs that are the most salient signals of free trade's negative consequences in the minds of those who filed the applications, as well as guard against some of the measurement concerns discussed by Kim and Pelc (2021). A full set of robustness checks can be found in the SI.

⁶Our results are substantively unchanged if we keep movers, although we drop them to avoid selection concerns if the decision / ability to move is a function of both attitudes and exposure.

⁷The robustness of these conclusions to differences in how we transform the exposure measure, and whether we drop certain types of petitions, is strongest for the decline in support for free trade. The positive coefficient on the increase in opposition is more sensitive, dropping below conventional thresholds for

Dependent Variables:	tradepolicy=="Favor"	tradepolicy=="Oppose"	tradepolicy=="DK"
Model:	(1)	(2)	(3)
<i>Variables</i>			
exposure	-0.0315** (0.0136)	0.0245** (0.0101)	0.0070 (0.0130)
<i>Fixed-effects</i>			
Respondent	Yes	Yes	Yes
Year	Yes	Yes	Yes
<i>Fit statistics</i>			
Observations	8,938	8,938	8,938
R ²	0.59499	0.55111	0.53864
Within R ²	0.00208	0.00206	0.00011

Clustered (Respondent) standard-errors in parentheses

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 1: Trade policy preferences predicted as a function of exposure to free trade’s negative consequences, operationalized as whether the respondent lived within 200 miles of any trade-related layoffs over the preceding 120 days prior to responding to the survey.

Even with the rich panel data, however, these results may capture spurious associations between place and attitudes that are unrelated to free trade’s negative consequences. Under the assumption that these place-based confounders are relatively time invariant (such as ruralness or local culture), we might find evidence of reverse causality in which 2011’s attitudes on trade are correlated with not just 2011’s measure of exposure but also to 2016’s measures of trade-related job loss. We reorganize our data to wide format and predict each wave’s attitudes as a function of either that wave’s exposure measure, or the exposure measures from the other two waves, controlling for a battery of respondent-level characteristics. We visualize these results in Figure 5, illustrating that the strongest associations are found when predicting attitudes on trade-related job losses over the four months prior to when the respondent participated in the survey. Interestingly, these results also indicate that the majority of the relationship is found in the 2011 and 2016 waves, with the pattern—if anything—reversed in 2020.

significance in some specifications, while the don’t know outcome increases to significant positive associations.

Placebo tests for reverse causality

Predicting attitudes by exposure measured at different periods




Figure 5: Placebo tests comparing strength of association between exposure and attitudes within a given survey wave (solid points) against leads and lags that break the temporal relationship (hollow points). X-axes indicate the year in which exposure is measured (aggregating over four months prior to the survey date), while panels indicate the year in which the attitude was measured. Gray squares indicate the correlation between exposure and favoring expanded free trade while red circles indicate the correlation between exposure and opposing expanded free trade.

One important question is whether these results are sensitive to how we define “exposure.” Insofar as we want to give the maximum support for the OEP framework, we search across a range of definitions of exposure across both space and time dimensions, aggregating layoffs by different distances and over different periods prior to when the respondent participated in the survey wave. Figure 6 summarizes these results, highlighting that the strongest evidence in favor of individuals updating their views on trade in response to increase trade-related layoffs is observed when we aggregate said layoffs over the preceding 4 to 9 months and over an area of between 100 and 200 miles. The figure also highlights that the bulk of the relationship is found not in increasing opposition to free trade per se, but rather in declining support for it, which appears to be evenly distributed between opposition

and uncertainty.

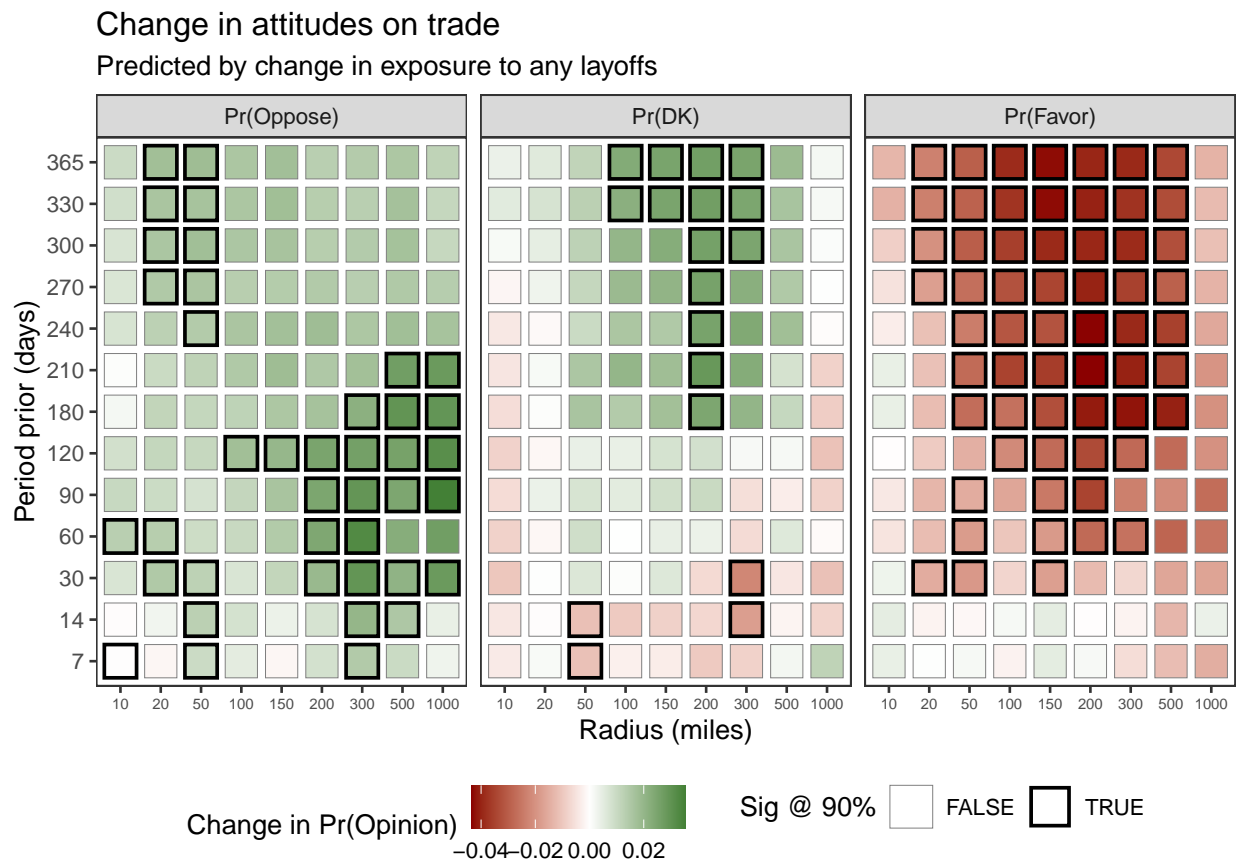


Figure 6: Each tile presents the $\hat{\beta}_1$ coefficient estimated from Equation 1, capturing the indicator for whether the respondent lived within x -miles from any trade-related layoffs (x -axis) over the preceding y -days (y -axis). Solid black borders indicate estimates that are significant at the 90% threshold.

4.2 Exposure’s effect on partisanship and vote choice

The preceding results provide descriptive evidence suggesting that attitudes on trade policy are correlated with proximity to trade-related layoffs in both space and time, with individuals living within approximately 200 miles from trade-related layoffs that occurred within 4 months prior to the survey indicating less approval for expanded trade policy. Do the same measures of exposure correlate with self-reported partisanship and vote choice? To inves-

tigate, we evaluate this hypothesis with a specification similar to that described above in Equation 1, except that we replace the trade opinion outcome with self-reported vote choice. As illustrated in Table 2, there is little evidence linking exposure to free trade’s negative consequences with actual vote choice, especially in the direction of increasing support for the Republican ticket between 2012, 2016, and 2020, at least when we use the same measure of exposure to free trade’s negative consequences that was most prognostic of increasing protectionism above (200 mile radius over the preceding four months). There is statistically significant evidence that exposure to trade-related layoffs increased political participation, although the substantive magnitude of this estimate (column 4) is quite small.

Dependent Variables: Model:	Democrat (1)	Republican (2)	Other Candidate (3)	No Vote (4)
<i>Variables</i>				
exposure	0.0047 (0.0072)	0.0003 (0.0069)	-0.0035 (0.0057)	-0.0015** (0.0006)
<i>Fixed-effects</i>				
Respondent	Yes	Yes	Yes	Yes
Year	Yes	Yes	Yes	Yes
<i>Fit statistics</i>				
Observations	8,964	8,964	8,964	8,964
R ²	0.89103	0.88447	0.53996	0.34945
Within R ²	0.00018	5.8×10^{-7}	0.00014	0.00052

Clustered (Respondent) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table 2: Vote choice predicted as a function of exposure to free trade’s negative consequences, operationalized as the logged applications aggregated over 200 miles and 120 days preceding when the respondent took the survey.

When we examine the robustness of these conclusions to different choices for the spatial and temporal “distances” over which we aggregate trade-related layoffs, we note suggestive evidence that exposure’s influence on vote choice is more proximate across both time and space. Specifically, Figure 7 indicates that layoffs which occurred within 150 miles in the preceding two months are correlated with increasing support for the Republican party and

declining support for the Democrats. Is it possible that individuals do in fact respond to trade-related layoffs by shifting their electoral support toward the increasingly protectionist Republican party, even if this occurs prior to an increase in their own protectionist attitudes? We turn to this question in the subsequent section.

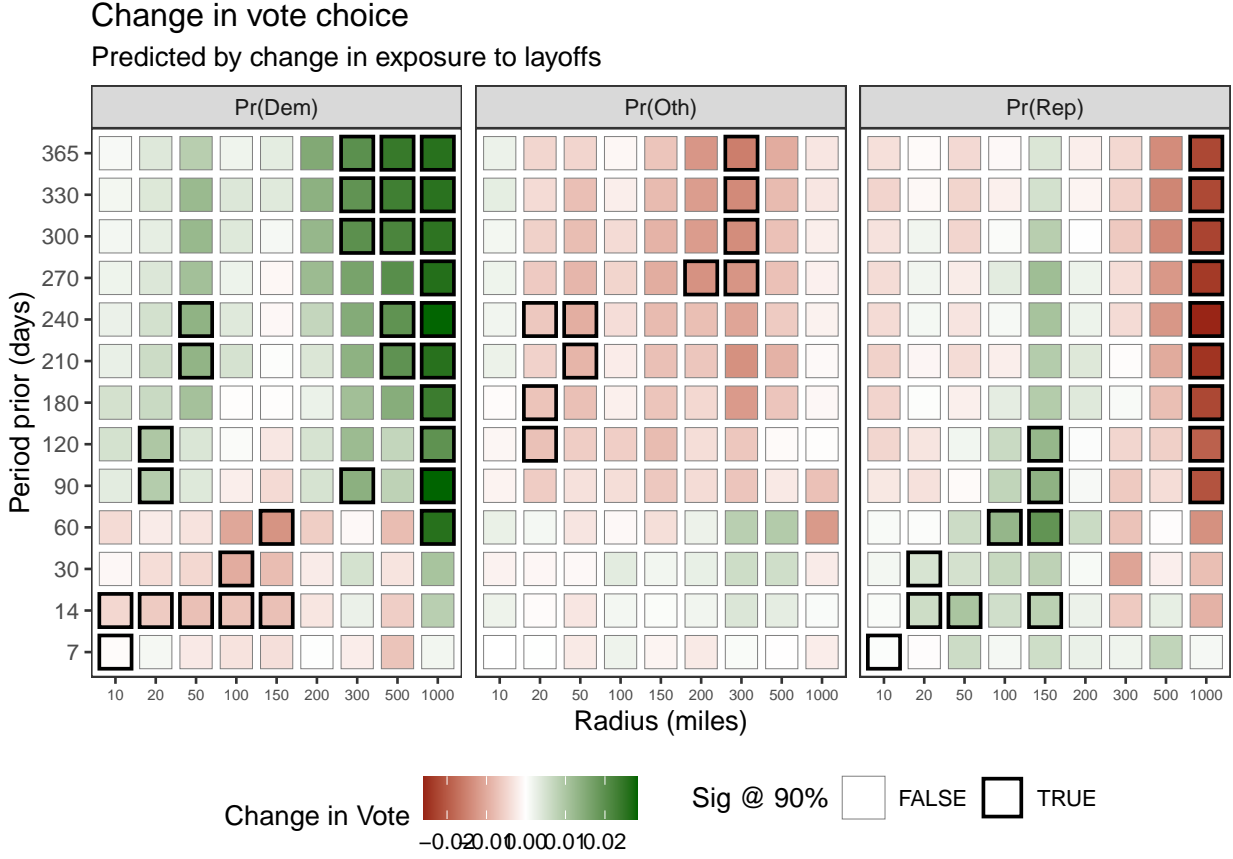


Figure 7: Each tile presents the $\hat{\beta}_1$ coefficient estimated from Equation 1, except where exposure to layoffs is used to predict variation in presidential vote, capturing the indicator for whether the respondent lived within x-miles from any trade-related layoffs (x-axis) over the preceding y-days (y-axis). Solid black borders indicate estimates that are significant at the 90% threshold.

4.3 Opinions on trade and presidential vote choice

To summarize our findings thus far, we find some evidence that attitudes on trade respond to exposure to trade-related layoffs, although these results are sensitive to the inclusion of

party ID controls and obtain only for the period between 2011 and 2016. Furthermore, we find no evidence indicating that the same exposure measures are prognostic of vote choice, although there is some evidence suggesting that vote choice is more influenced by greater proximate exposure, defined in terms of both time and space.

Here, we investigate whether trade attitudes themselves are correlated with vote choice. To do so, we run two descriptive regressions that directly capture the intuition of the OEP versus PMR debate. Our first treats the vote as the outcome and regresses it on the respondents' protectionist attitudes, measured in 2011. Our second treats protectionist attitudes as the outcome and regresses it on whether the respondents supported the Republican candidate in 2012. Where possible, we control for lagged values of the dependent variable, although the results are robust to its omission.

Figure 8 suggests a simple but powerful resolution to the OEP versus PMR debate: attitudes on free trade—measured in 2011—are never significant predictors of GOP support in the 2012, 2016, or 2020 elections and are decreasingly prognostic of vote choice. Conversely, these attitudes are strikingly well explained by the respondents' support for the Republican candidate in 2011 as soon as former President Trump campaigned aggressively on a protectionist platform. Put plainly, this evidence suggests that Republicans shifted their stance on free trade to align with former President Trump's views, rather than selecting a presidential candidate based on their own trade beliefs.

However, the preceding results—by pinning both predictors to their values recorded in the first survey wave—suppress subsequent variation that might reveal a more nuanced perspective on the relative explanatory power of OEP versus PMR. Figure 9 displays changes in attitudes on trade among consistent partisans (top row) and changes in vote choice among those with consistent trade policy preferences.⁸

⁸We define “consistent” to mean those for whom vote choice and attitudes did not change over the period of analysis.




Figure 8: The left panel shows that voting in 2012, 2016, and 2020 is increasingly poorly predicted by one’s opposition to free trade measured in 2011. The right panel suggests that one’s opposition to free trade in 2011, 2016, and 2020 is increasingly correlated with one’s intended support for the Republican party.

In the top row of plots, we find striking evidence of partisans changing their attitudes between the first and last waves of the survey, particularly among consistent Democrats who clearly updated their attitudes on trade to reflect the growing partisan divide on the issue. Conversely, the bottom pair of plots examines whether respondents with consistent attitudes on trade updated their vote choice between 2011 and 2020. It finds much milder evidence of voting on the basis of policy alignment. While 11% of the pro-free trade Republicans in 2011 shifted support to the Democrats in 2020, and 22% of protectionist Democrats in 2011 switched sides in 2020, these movements—consistent with the rational actor underpinning the OEP framework—are marginal when compared to the flexibility of policy preferences.




Figure 9: Changes in attitudes on trade among consistent partisans are more dynamic (top) than changes in vote choice among those with consistent trade policy preferences (bottom).

Descriptively, it seems that voters die with their partisan, not protectionist, boots on.

These patterns are consistent with the literature on political behavior and partisanship as a social identity, which suggests that voters rarely vote on policies within any degree of consistency in even the best of circumstances (Campbell et al., 1960; Achen and Bartels, 2017; Green, Palmquist and Schickler, 2004). Recent scholarly attention to the psychology of international political economy urges extreme caution when attempting to gauge the consistency of American trade attitudes with their proclaimed partisanship (Mutz, 2021). That is, there is very little—if any—consistency there, particularly when considering a voter’s ability to match free trade policy positions with candidates (spoiler: most can’t) (Mutz, 2021). Despite the abstract nature of international trade and its link to the economy, which presents challenges to even the most educated of voters, the us-versus-them and competitive (versus cooperative) social dominance paradigms might be what drive protectionist attitudes more generally (Mutz, 2021). What, then, do our regression results suggest about the Republican-leaning protectionist trajectory before the election of President Trump in 2016, and the continuation of that trend until 2020?

It is tempting to draw the conclusion that protectionist attitudes *are* more prevalent in the Republican camp in the last decade minus, and they *are* ultimately increasing the likelihood of a vote for the Republican presidential candidate. Figure 10 shows a considerable effect with an issue coefficient of 0.174 (standard error 0.014). Those who favor the OEP framework might conclude that trade protectionist attitudes do result in a greater likelihood of Republican presidential votes—and with perhaps a more substantial effect than once thought possible. Indeed, it might surprise many that trade attitudes influence presidential voting at all, as we have previously noted. However, there are challenges of internal validity that must be reconciled. Specifically, it is not possible to tell whether the protectionist inclination among our respondents is actually motivating their voting for Republican candidates, or if they are simply deciding upon the candidates they like *first* and adopting the

candidates' trade policy positions as their own *later*. We attempt to reconcile such concerns with the two-wave, time-lagged model. Using this methodology, we quickly recognize that the conventional test might not provide the full story. When we convert the conventional test into the time-lagged specification that controls for priors, the issue weight decreases substantially, and we also find evidence of priming, meaning our respondents likely did not arrive at their trade preferences independently but instead had the nudge of elite cuing.




Figure 10: Conventional Test Marginal Effects of Protectionist Trade Preferences on Republican Presidential Vote (2012-2020)

The two-waved test results in Table 3 are further evidence in favor of the PMR paradigm for three reasons. First, the weight that individuals place on trade policy when casting a vote for the Republican candidate is considerably dampened from that displayed in Figure 10. That is, the t-statistic for the 2012 trade issue weight is a drop of more than two-thirds from the issue weight depicted by the conventional test. Second, the weight placed

on trade policy does increase between 2012 (0.05) and 2020 (0.09), which provides evidence of priming and may account for the raised in salience of protectionism as a Republican platform in the run-up to the 2016 election and afterwards. Third, party identification absorbs a substantial proportion of explanatory power in the outcome. While the conventional test might suggest that former President Trump raised the salience of protectionism in an outsized way in 2016 and with lasting effects, the time-lagged results motivate a qualification of that conclusion. Trade preferences—as a mark of independent, baseline convictions among individuals—have proved susceptible to elite priming over last decade. It is especially difficult to conclude that any strength of conviction on trade that does exist among individuals is not immediately swallowed by overwhelming explanatory power of partisan identification.

Table 3: Effects of Protectionist Trade Preferences on Republican Presidential Vote Choice

	<i>Dependent variable:</i>	
	GOP Vote 2012	GOP Vote 2020
	(1)	(2)
Trade Preference 2011	0.05*** (0.02)	0.09*** (0.02)
PID 2011	0.95*** (0.02)	0.89*** (0.02)
Constant	−0.04*** (0.01)	0.04*** (0.01)
Observations	3,235	3,235
R ²	0.50	0.43
Residual Std. Error (df = 3232)	0.34	0.38
<i>Note:</i>	*p<0.1; **p<0.05; ***p<0.01	

5 Discussion and Conclusion

This research investigates the relationship between exposure to trade-related layoffs and voting behavior, with specific attention to the Republican party and attitudes towards free trade from 2011 to 2020. Employing a robust methodology, including conventional and time-lagged tests, our findings reveal a minimal direct connection between trade exposure and vote choice, with a more pronounced shift in attitudes away from free trade, especially among Republican voters. The findings underscore the significant role of partisan identification, which substantially influences voting behavior and attitudes towards trade, often overshadowing individual trade preferences. We further document empirical patterns that are consistent with elite priming, suggesting that individuals' trade preferences have grown increasingly susceptible to political influence over the last decade, corresponding to the politicization of trade over the same period.

These results highlight the growing politicization of the issue of free trade, and the erosion of the microfoundations of the Open Economy Politics (OEP) framework. Contrary to the OEP framework, which posits a direct and significant relationship between economic exposure and political preferences, our findings highlight a minimal and decreasing impact of trade exposure on voting behavior. Instead, we document a pronounced and growing influence of Partisan Motivated Reasoning (PMR) on individuals' attitudes towards free trade and voting choices. Our study reveals that partisan identification substantially overshadows and mediates the relationship between trade preferences and voting behavior, indicating a shift towards increased PMR over the period of analysis. The evidence of elite priming further underscores this trend, demonstrating the susceptibility of individuals' trade preferences to political influence, thereby reinforcing the ascendancy of PMR in shaping political attitudes and behaviors in the context of free trade issues.

These findings suggest three gaps in our understanding that can be addressed with

future research. First, to what extent was the rational actor model undergirding OEP better able to explain beliefs and voting behavior prior to the period analyzed here? While trade’s increased salience during the 2016 election and former President Trump’s time in office makes it more plausible that voters evaluated candidates on the basis of trade, the logic cuts both ways. Bullock (2020) highlights the importance of “distinctiveness” of party positions as a crucial scope condition for motivated reasoning, and our period of analysis certainly fits the bill as one in which Democrats and Republicans became far more distinctive on the issue of trade policy. While we know that aggregate patterns in voting behavior correlate with the geographic distribution of free trade’s winners and losers dating back to the 1990s (see Autor, Dorn and Hanson 2013; Autor et al. 2016, 2017; Jensen, Quinn and Weymouth 2017; Baccini and Weymouth 2021), and that cross-sectional surveys find similar patterns between trade attitudes and exposure (see Dippel, Gold and Heblich 2015; Colantone and Stanig 2018*a,b*), a more careful test using individual-level panel data from earlier periods is required to confirm the microfoundations of the OEP framework.

Second, assuming that these microfoundations are supported in earlier periods, why didn’t political elites exploit this untapped reservoir of political behavior earlier? At least part of the increasing correspondence between trade attitudes and Republican support between 2012 and 2016 is from former President Trump’s decision to center his campaign on trade. Was he able to do so thanks to some change in the underlying political environment, or was it because of his norm-breaking status as a Republican on other dimensions that allowed him to adopt a protectionist stance at odds with the broader GOP platform? Conversely, if individual-level panel data evidence from less politicized periods still fails to support the self-interested intuition of OEP, how then can we explain the predictive power of the models on which it is built? And what then explains the stronger descriptive patterns linking regional variation in free trade’s negative consequences to protectionism and support for populist parties? Our results suggest that the answer might be found in organized interests, localized elite cues, or perhaps “coalition merchants” (Noel, 2012), and—like similar

conclusions drawn by Rho and Tomz (2017)— highlights the importance of incorporating other theoretical traditions into the OEP understanding of how interests and institutions interact to shape foreign policy.

Third and finally, the patterns documented here carry long-term implications for global trade relations and policy negotiations. Are countries moving towards more protectionist or free trade policies as a result of these internal political shifts? How does the erosion of the open economy politics (OEP) framework, as evidenced in the U.S., resonate in other global contexts, and what does it mean for international cooperation and conflict in trade? Exploring these questions is essential for anticipating and navigating the future landscape of global trade, as internal political dynamics continue to evolve and intersect with international economic interactions and policies.

Appendices

A Construct validity of TAA applications

Each TAA application contains three pieces of information that combine to capture the core quantity of interest for our respondents: exposure to free trade’s negative consequences. First, these applications include the physical address of the location at which workers were laid off due to some aspect of free trade. We process every address via Google’s geocode package for R and calculate the latitude and longitude of the firm, store, plant, or factory to a very precise degree of accuracy.⁹ This precision allows us to then assign each respondent’s exposure to free trade’s negative consequences as a function of their physical proximity to the location where the layoffs occurred.

We examine a variety of methods for assigning respondents to layoffs, including adding up the total number of applications within an X-mile radius from the respondent, or via a distance-weighted sum of all applications. However, while the precision of our latitude and longitude coordinates for the TAA data is precise down to the physical address, our precision for the survey respondents is substantially worse. While zip codes are some of the smallest geographic units available in the United States (there are roughly 41,704 zip codes in the United States in 2023), they are nevertheless of far lower resolution than the physical addresses found in the TAA data. Furthermore, they are not strictly geographic units at all, but rather a collection of mail delivery routes associated with different post offices. The largest zip code covers an area of almost 10,000 square miles between Las Vegas and Reno, while the smallest is in Manhattan and covers barely two city blocks. Given this variation, we

⁹The resolution is equivalent to Google’s Street View, or roughly 11 square meters. However, we reduce this resolution to 111 square meters for computational efficiency, given that our respondents themselves are geolocated with substantially more noise.

are clearly introducing measurement error when we assign different respondents to different levels of exposure based on the geographic distance between a precisely located factory, and a noisily located individual. We provide a more thorough discussion of this empirical challenge in our Supporting Information, but emphasize here that this is (we believe) the best we can do with the available data.

The second feature of the TAA data is that it contains an estimate of the total number of workers who were laid off due to free trade *in the minds of those affected*. A persistent challenge with operationalizing free trade's negative consequences is that we are unsure of whether an individual working in a given industry or possessing a certain skill that is threatened by import competition can attribute their welfare outcomes to trade specifically. Indeed, the difficulty among leading experts in the fields of economics and political science to reach a consensus on which characteristics actually matter for identifying trade's winners and losers underscores just how non-trivial the attribution task is for the average voter. (Of course, this very difficulty supports our broader claim that the individual-level pathway linking trade's consequences to politics is unlikely to be determinative.)

Nevertheless, there are two important caveats to our interpretation of TAA applications as capturing layoffs that are attributed to trade in the minds of individuals. First, obviously, the individuals who submit TAA applications are unlikely to be the same individuals who were empaneled in our survey. Indeed, we know that not all TAA applications are filed by the affected workers themselves. Some are initiated by unions, others by state officials representing the Department of Labor, still others by the firms who laid off the workers in the first place. And even among the applications which were filed by the directly affected workers, even if by some vanishingly small chance these workers also happened to be the respondents in our survey, we have no way of linking them.

Second, and more concerningly, there is evidence suggesting that the decision to apply for TAA support is endogenous to politics, both in the sense that the applicants themselves

might be swayed by ideology or partisanship, and also in the sense that their local politicians hold differing views on the value and appropriateness of programs like the TAA, and might differentially advertise or suppress its availability.

As above, we acknowledge both of these limitations, but highlight that our panel data allows us to rule out the most concerning selection effects suggested by these caveats. If individuals living in more pro-Republican areas are less likely to file for TAA, we are nevertheless able to hold this constant via the panel data. Under the reasonable assumption that ideology or partisanship are time-invariant qualities of our respondents (an assumption we support with descriptive results in our Supporting Information), implementing respondent-fixed effects ensures that we are identifying variation in attitudes and voting behavior that is associated with the change in exposure, even if measured with error.

The final feature of the TAA data that bears mentioning is the determination itself. Every application is reviewed by an employee at the Department of Labor who conducts a formal investigation of the claims put forth in the application. This involves confirming that the putative reason for the layoffs were indeed due to one of either import competition (determination code C1 in the data), shifts in production (C2), downstream consequences of either of the preceding (CD), or upstream consequences of either of the preceding (CU). In addition, an application might be denied (D), partially certified (P), or terminated (T). Our data contains all of these codes, allowing us to more carefully parse which reasons for the layoffs are most prognostic of attitudes and behaviors. (Importantly, one might wonder whether denials should be less prognostic of attitudes since they reflect layoffs that were not actually due to free trade, or more prognostic of attitudes since they reflect layoffs thought to be due to free trade, but which did not receive the additional support expected by the applicants and affected workers.)

B Variable Importance

A potential concern with the horse race regression presented in the main results is that they might be sensitive to specification choices. For example, is it appropriate to model these effects additively in a linear regression specification? And even with the sensitivity analyses which evaluate the strength of the results to both geographic and temporal distance, are we neglecting crucial but nuanced variation in the historic exposure to layoffs among our respondents?

To evaluate, we adopt an alternative measure of exposure in which we calculate a distance-weighted measure of exposure to layoffs using a Γ distribution with a monotonic decay over distance, and calculate the per quarter exposure for each respondent, dating back to 2000. This generates a k length vector of exposure measures for each respondent, where k is equal to the number of quarters between Q1 2000 and Q3 of either 2011, 2016, or 2020 when the survey response was recorded. We visualize an example of these vectors in Figure 11, where each line is a survey respondent’s exposure to trade-related layoffs, and the x-axis is the quarter of interest.

Spatially, the Bayesian framework might suggest that layoffs further away are less salient (i.e., $\frac{\partial \sigma_{\text{layoffs}}^2}{\partial \text{distance}} < 0$), prompting us to use a distance-weighted measure with a decay rate such that layoffs further than 300 miles away are effectively given zero weight.¹⁰ However, temporally the intuition is less clear. On the one hand, it is sensible to expect more recent layoffs are more influential of preferences. But it is also reasonable to expect that one’s historic exposure to trade-related job loss is also prognostic of their attitudes, and might condition the influence of more recent exposures. For example, an individual who is consistently exposed to layoffs might be less influenced by a round of job losses in the quarter preceding the survey than an individual who was never exposed until the same quarter. Yet

¹⁰For a description of this method, please see the SI.




Figure 11: Exposure histories for a random sample of respondents.

we might also expect that the consistent and repeated exposures over time has also made the former individual hold more negative attitudes, which a focus on only the most recent quarter would miss.

Determining the precise functional form in which these vectors combine to influence attitudes in a linear regression context is heroic enough to be a fool’s errand. Instead, we turn to machine learning methods – specifically random forests – to non-parametrically model attitudes on trade as a function of these k -length vectors of exposure, as well as partisanship. We run these models separately for 2011, 2016, and 2020, and evaluate the research question using permutation tests of variable importance. These tests estimate the random forest on the full data and record the predictive accuracy of the observed data.

Then we randomly reshuffle predictors of interest, breaking the empirical relationship and re-record the predictive accuracy. The measure of variable importance is the proportional reduction in model fit associated with permuting each predictor.

We visualize the findings in Figure 12, highlighting the declining predictive power of trade-related layoffs between 2011 and 2020, which is paired with the increasing predictive power of party ID over the same period. These results further support our conclusion that the politicization of free trade between 2011 and 2020 has eroded the microfoundations of the OEP framework, and tilted this issue domain further toward partisan motivated reasoning.




Figure 12: Variable importance (y-axis) for layoff vectors (white bars) and partisanship (solid bars) by survey wave (x-axis).

References

- Abramowitz, Alan I and Steven Webster. 2016. “The rise of negative partisanship and the nationalization of US elections in the 21st century.” *Electoral Studies* 41:12–22.
- Achen, Christopher H. 1992. “Social psychology, demographic variables, and linear regression: Breaking the iron triangle in voting research.” *Political behavior* 14:195–211.
- Achen, Christopher H and Larry M Bartels. 2017. Democracy for realists. In *Democracy for Realists*. Princeton University Press.
- Autor, David, David Dorn, Gordon Hanson and Kaveh Majlesi. 2017. “A note on the effect of rising trade exposure on the 2016 presidential election.” *Appendix to “Importing Political Polarization* pp. 3139–83.
- Autor, David, David Dorn, Gordon Hanson and Kaveh Majlesi. 2020. “Importing political polarization? The electoral consequences of rising trade exposure.” *American Economic Review* 110(10):3139–3183.
- Autor, David, David Dorn, Gordon Hanson, Kaveh Majlesi et al. 2016. “Importing political polarization.” *Massachusetts Institute of Technology Manuscript* .
- Autor, David H, David Dorn and Gordon H Hanson. 2013. “The China syndrome: Local labor market effects of import competition in the United States.” *American economic review* 103(6):2121–2168.
- Baccini, Leonardo and Stephen Weymouth. 2021. “Gone for good: Deindustrialization, white voter backlash, and US presidential voting.” *American Political Science Review* 115(2):550–567.
- Barone, Guglielmo and Helena Kreuter. 2021. “Low-wage import competition and populist backlash: The case of Italy.” *European Journal of Political Economy* 67:101970.

- Bartels, Larry M. 1993. "Messages received: The political impact of media exposure." *American political science review* 87(2):267–285.
- Bisbee, James and B Peter Rosendorff. 2020. "Shocking the Vulnerable: Job Insecurity, Local Sociotropism and Anti-Globalization Sentiment." *International Political Economy Society Working Paper*, URL .
- Bullock, John G. 2020. *Party cues*. Oxford University Press, New York, NY.
- Campbell, Angus, Philip E Converse, Warren E Miller and Donald E Stokes. 1960. *The American voter*. University of Chicago Press.
- Colantone, Italo and Piero Stanig. 2018a. "Global competition and Brexit." *American political science review* 112(2):201–218.
- Colantone, Italo and Piero Stanig. 2018b. "The trade origins of economic nationalism: Import competition and voting behavior in Western Europe." *American Journal of Political Science* 62(4):936–953.
- Deng, Liuchun, Pravin Krishna, Mine Zeynep Senses and Jens Stegmaier. 2022. Trade, Human Capital, and Income Risk. Technical report National Bureau of Economic Research.
- Dippel, Christian, Robert Gold and Stephan Heblich. 2015. Globalization and its (dis-) content: Trade shocks and voting behavior. Technical report National Bureau of Economic Research.
- Druckman, James N and Mary C McGrath. 2019. "The evidence for motivated reasoning in climate change preference formation." *Nature Climate Change* 9(2):111–119.
- Feler, Leo and Mine Z Senses. 2017. "Trade shocks and the provision of local public goods." *American Economic Journal: Economic Policy* 9(4):101–143.
- Fiorina, Morris P. 1981. "Retrospective voting in American national elections." (*No Title*).

- Green, Donald P, Bradley Palmquist and Eric Schickler. 2004. *Partisan hearts and minds: Political parties and the social identities of voters*. Yale University Press.
- Guisinger, Alexandra. 2016. "Information, gender, and differences in individual preferences for trade." *Journal of Women, Politics & Policy* 37(4):538–561.
- Hainmueller, Jens and Michael J Hiscox. 2006. "Learning to love globalization: Education and individual attitudes toward international trade." *International Organization* 60(2):469–498.
- Highton, Benjamin and Cindy D Kam. 2011. "The long-term dynamics of partisanship and issue orientations." *The Journal of Politics* 73(1):202–215.
- Hiscox, Michael J. 2001. "Class versus industry cleavages: inter-industry factor mobility and the politics of trade." *International Organization* 55(1):1–46.
- Inglehart, Ronald F and Pippa Norris. 2016. "Trump, Brexit, and the rise of populism: Economic have-nots and cultural backlash."
- Jensen, J Bradford, Dennis P Quinn and Stephen Weymouth. 2017. "Winners and losers in international trade: The effects on US presidential voting." *International Organization* 71(3):423–457.
- Karol, David. 2009. *Party position change in American politics: Coalition management*. Cambridge University Press.
- Kim, Sung Eun and Krzysztof J Pelc. 2021. "Trade Competition and Worker Compensation: Why Do Some Receive More than Others?" *International Studies Quarterly* 65(1):109–121.
- Lake, David A. 2009. "Open economy politics: A critical review." *The Review of International Organizations* 4:219–244.

- Lenz, Gabriel S. 2013. *Follow the leader?: how voters respond to politicians' policies and performance*. University of Chicago Press.
- Malgouyres, Clément. 2017. "Trade shocks and far-right voting: Evidence from French presidential elections." *Robert Schuman Centre for Advanced Studies Research Paper No. RSCAS 21*.
- Mansfield, Edward D and Diana C Mutz. 2009. "Support for free trade: Self-interest, sociotropic politics, and out-group anxiety." *International organization* 63(3):425–457.
- Margalit, Yotam. 2011. "Costly jobs: Trade-related layoffs, government compensation, and voting in US elections." *American Political Science Review* 105(1):166–188.
- Margalit, Yotam. 2012. "Lost in globalization: International economic integration and the sources of popular discontent." *International Studies Quarterly* 56(3):484–500.
- Mutz, Diana C. 2021. *Winners and losers: The psychology of foreign trade*. Vol. 27 Princeton University Press.
- Noel, Hans. 2012. "The coalition merchants: The ideological roots of the civil rights realignment." *The Journal of Politics* 74(1):156–173.
- Pierce, Justin R and Peter K Schott. 2016. "The surprisingly swift decline of US manufacturing employment." *American Economic Review* 106(7):1632–1662.
- Rho, Sungmin and Michael Tomz. 2017. "Why don't trade preferences reflect economic self-interest?" *International Organization* 71(S1):S85–S108.
- Ricardo, David. 1821. *On the principles of political economy*. J. Murray London.
- Rickard, Stephanie J. 2020. "Economic geography, politics, and policy." *Annual Review of Political Science* 23:187–202.

- Scheve, Kenneth, David Stasavage et al. 2006. "Religion and preferences for social insurance." *Quarterly Journal of Political Science* 1(3):255–286.
- Scheve, Kenneth F and Matthew J Slaughter. 2001. "What determines individual trade-policy preferences?" *Journal of International Economics* 54(2):267–292.
- Sides, John, Chris Tausanovitch and Lynn Vavreck. 2022. *The bitter end: The 2020 presidential campaign and the challenge to American democracy*. Princeton University Press.
- Stolper, Wolfgang F and Paul A Samuelson. 1941. "Protection and real wages." *The Review of Economic Studies* 9(1):58–73.
- Vavreck, Lynn. 2009. *The message matters: The economy and presidential campaigns*. Princeton University Press.
- Zechman, Martin J. 1979. "Dynamic models of the voter's decision calculus: Incorporating retrospective considerations into rational-choice models of individual voting behavior." *Public Choice* 34(3-4):297–315.