

Can Pigou at the Polls Stop Us Melting the Poles?

Soren Anderson, Ioana Marinescu, Boris Shor

Abstract: Economists recommend Pigouvian taxes as the most efficient way to fight climate change. Yet, carbon taxes are difficult to implement politically. To understand why, we study Washington State's two failed carbon tax referendums from 2016 and 2018—the first such votes in the United States. We find that average voters' opposition to the carbon tax can partly be explained by the anticipation of higher energy costs. Meanwhile, ideology—as measured by voting on other initiatives—explains 90% of variation in voting across precincts. These results suggest that ideology plays a crucial role in driving opposition to carbon taxes. We find that revenue recycling interacts with ideology: conservatives preferred the 2016 revenue-neutral policy, while liberals preferred the 2018 green-spending policy. Finally, we forecast that no other state is liberal enough to pass Washington's policies. Thus, opinion surveys showing majority support for the carbon tax can be misleading.

JEL Codes: D72, H23, H71, H72, Q52, Q54, Q58

Keywords: Pigouvian taxes, carbon taxes, voting behavior, willingness to pay, state environmental policy

Soren Anderson is at Michigan State University and the National Bureau of Economic Research (NBER) (sta@msu.edu). Ioana Marinescu is at the University of Pennsylvania and NBER (ioma@upenn.edu). Boris Shor is at the University of Houston (bshor@uh.edu). We thank Aria Kovalovich, Doug Murdoch, Anna Terkelsen, and Yuxian Xiao for valuable research assistance; Nicholas Pharris at the Washington Secretary of State for providing data and GIS expertise; Yoram Bauman, Gail Gatton, Ben Silesky, and Greg Small for helpful background and detailed comments; Matt Kahn, Jim Hines, Matthew Incantalupo, and Sheila Olmstead (the editor) and three anonymous referees for detailed comments; and seminar participants at University of California Davis, Energy Policy Institute at Chicago, University of California Berkeley Energy Camp, Michigan State University, University of Michigan, and Western Michigan University Energy and Environmental Economics Day Workshop, Purdue University, Heartland Environmental and Resource Economics Workshop at Illinois, Oberlin College, the Allied Social Science Associations annual meetings in Philadelphia, Michigan State University, the Midwest Political Science Database data: <https://doi.org/10.7910/DVN/GORATE>

Received April 8, 2020; Accepted October 7, 2022; Published online May 10, 2023.

Journal of the Association of Environmental and Resource Economists, volume 10, number 4, July 2023.

© 2023 The Association of Environmental and Resource Economists. All rights reserved. Published by The University of Chicago Press for The Association of Environmental and Resource Economists.
<https://doi.org/10.1086/722970>

ESTIMATES FOR 2020 PUT THE SOCIAL COST OF CARBON at \$14–\$152 per metric ton of CO₂.¹ A Pigouvian tax on emissions can internalize the social cost of carbon—as can a system of tradable emissions permits (“cap and trade”). Price incentives equalize the marginal cost of reducing emissions across different sources and thereby minimize the total cost of meeting a given emissions target. Meanwhile, revenue from a carbon tax can be used (“recycled”) to lower other taxes or to mitigate impacts on low-income households and energy-intensive industries. This logic appeals to economists from across the political spectrum, as well as to many commentators and politicians, as exemplified by the informal Pigou Club and more formal Climate Leadership Council.² Yet these economic arguments inevitably confront a key political hurdle: carbon taxes lead to higher energy prices, which are highly unpopular (see Knittel 2014). How then can a carbon tax be enacted in a democracy? One view is that forming a bipartisan coalition would be easier if carbon taxes were more acceptable to conservatives who tend to favor limited government. This view supports using carbon tax revenue to lower other taxes in a revenue-neutral manner. An alternative view sees the bipartisan approach as doomed to failure—presuming that conservatives would reject a carbon tax of any flavor, while progressives would actually prefer to increase government spending. This alternative view supports spending carbon tax revenue on green projects and social programs to maximize liberal support.

How important are energy costs and political considerations in explaining voters’ willingness to pay for carbon taxes? Does revenue recycling matter? Surveys indicate that support is largely driven by political ideology (Egan and Mullin 2017) and can depend on how the revenue is spent (Amdur et al. 2014; Kotchen et al. 2017). Do these result hold up in real-world voting? To answer these questions, we study voting patterns in two recent, failed carbon tax initiatives in Washington State. Initiative 732

Association meeting in Chicago, the State Politics and Policy Conference meeting in College Park, the NBER Summer Institute, Macalester College, Brigham Young University, and Stanford University for many helpful comments and suggestions. Ioana Marinescu acknowledges financial support from the Kleinman Center for Energy Policy at the University of Pennsylvania.

1. See https://www.whitehouse.gov/wp-content/uploads/2021/02/TechnicalSupportDocument_SocialCostofCarbonMethaneNitrousOxide.pdf. The social cost of carbon is the present-discounted value of the stream of present and future economic damages from climate change, caused by releasing one additional ton of carbon dioxide into the atmosphere. Also see <https://carbonpricingdashboard.worldbank.org/>.

2. The “Pigou Club” is a designation credited to founding member and conservative economist Greg Mankiw and includes fellow economists such as Paul Krugman, Ed Glaeser, Kevin Hassett, and William Nordhaus; Democratic politicians such as Al Gore; Republican politicians such as Rex Tillerson and Lindsey Graham; conservative/libertarian pundits such as David Frum and Megan McArdle; liberal/progressive pundits such as Bill Nye and David Leonhardt; and many others from across the political spectrum. The Climate Leadership Council is a bipartisan policy institute founded by Ted Halstead in 2017, which counts former Federal Reserve chairs Ben Bernanke and Janet Yellen among its founding members.

(I-732) from 2016—the first carbon tax referendum in the United States—would have lowered the state sales tax from 6.5% to 5.5% and matched the federal Earned Income Tax Credit (EITC) by 25%. Initiative 1632 (I-1631) from 2018 would have devoted 95% of carbon tax revenue to green projects like renewable energy. I-732 and I-1631 were otherwise quite similar, with tax rates starting at \$15 per ton of carbon dioxide (tCO_2) and rising gradually thereafter (see table 1). Thus, these initiatives offer an unprecedented opportunity to measure voter preferences for two substantively different

Table 1. Comparing Two Flavors of a Carbon Tax: I-732 versus I-1631

	I-732	I-1631
Year	2016	2018
Provisions	Revenue-neutral carbon tax swap \$15/ tCO_2 in July 2017, \$25/ tCO_2 in July 2018, then increase 3.5% per year to \$100/ tCO_2 Slower phase-in for farmers and public transportation Reduce state sales tax by 1% from 6.5% to 5.5% Reduce state business and occupa- tion tax on manufacturing busi- nesses to .001% Offer working families tax rebate (25% match on federal Earned Income Tax Credit)	Carbon emissions fee and spending \$15/ tCO_2 in January 2020, then increase \$2/ tCO_2 per year until state's emissions reduction goals met Levied on "large emitters" using and distributing fossil fuels Revenue to three funds: (1) 70% air quality and energy projects, (2) 25% water quality and forest projects, and (3) 5% for communities Establish public oversight board to determine spending from funds Create three panels to make spending recommendations to public oversight board
Results	40.75% Yes, 59.25% No	43.44% Yes, 56.56% No
Spending for	\$3.2 million	\$16.4 million
Spending against	\$1.4 million	\$31.6 million
Top spenders for	Peter Kelly (\$125,000)	Nature Conservancy (\$3.4 mil- lion), League of Conservation Voters (\$1.4 million), Bill Gates and Michael Bloomberg (\$1 million each)
Top spenders against	Kaiser Aluminum (\$450,000)	BP America (\$13.15 million), Phillips 66 (\$7.2 million), Andeavor (\$6.1 million)

Source. Ballotpedia.

versions of a carbon tax based on actual voting. Our analysis relies on detailed, precinct-level elections data from Washington State for 2016 and 2018.

First, we investigate what drives support for the carbon tax among Washington's voters, focusing on ideology and WTP for the carbon tax. To do so, we use observed variation in voter characteristics across precincts and variation in carbon tax policy parameters between 2016 and 2018. We seek to answer two questions. One: What explains the wide variation in support for the carbon tax across precincts? Overwhelmingly, we find the answer to be ideology. We infer latent ideology from votes on a dozen other policies, such as clean car subsidies and a higher minimum wage.³ Ideology alone explains 91% of the variation in vote shares across precincts and predicts votes better than partisanship (as measured by the Republican vote share). In contrast, demographic variables—including key proxies for carbon tax incidence, such as vehicle ownership and home size—predict relatively little of the variation across precincts. Ideology also explains changes in support between 2016 and 2018. I-732 was designed by an economist to be revenue neutral in an explicit appeal to political moderates, while I-1631 was designed by a progressive coalition with social and environmental justice objectives. Consistent with these different strategies, we find that I-1631 performed better in liberal precincts, where it picked up 3.3 percentage points relative to I-732—but worse in conservative precincts, where it lost 0.65 percentage points relative to I-732. Two: Why did the carbon tax fail to gain majority support? Here, we find that part of the answer is opposition to higher energy prices. We calculate the direct tax incidence due to personal energy consumption for each precinct and measure its impact on voting, which allows us to back out the distribution of WTP for the overall policy across precincts. For the average precinct, we find that the tax incidence (about \$100 per person per year) is important in driving opposition to the carbon tax. Yet, this tax incidence is similar across precincts, which explains why ideology rather than pocketbook issues accounts for most of the variation in vote shares across precincts.

Second, we use our regression results to predict how the vote in Washington would have changed were Washington assigned the same ideology and demographics as other states. This exercise is akin to an out-of-sample forecast for the carbon tax vote shares in other states. Based on this exercise, we find that the best chance for Pigou at the polls would be a carbon tax modeled on I-1631 in Massachusetts (49.1% support). We forecast that this policy could have the highest possible chance of passing in Vermont (50.6% support), but the Green Mountain State lacks the popular initiative mechanism. Meanwhile, the best chance for a revenue-neutral policy like I-732 is in California (47.1% support). Again, I-732 would do better in both New York and Hawaii (47.2%

3. We treat ideology as a latent concept that we measure using the observable signals of support for all ballot initiatives during the time of the relevant election. This is a technique used in a variety of studies (Snyder 1996; Gerber and Lewis 2004; Masket and Noel 2012). These studies validate the notion that district ideological opinion is mostly one dimensional.

and 47.8%), but neither state has the popular initiative. The states that show the highest predicted support generally already have some form of carbon pricing, typically through cap and trade. In fact, Washington state passed its own cap-and-trade legislation in May 2021. Passing a carbon tax via popular referendum appears to be more difficult than achieving carbon pricing through ordinary legislation.

This study contributes to the literature in political economy and public economics that uses votes in referendums to measure policy preferences. We are the first to analyze voting on a carbon tax in the United States, where climate issues are highly politicized. A carbon tax is arguably the most salient policy in environmental economics research: climate change looms as the largest environmental problem we face, carbon taxes are almost always the benchmark by which other policies (e.g., fuel economy standards) are judged, and Pigouvian taxes conceptually underpin the entire subfield of environmental policy design (e.g., via their symmetry to cap-and-trade programs). The closest papers are Bornstein and Lanz (2008), who study voting on three proposals to tax the energy (but not carbon) content of nonrenewable energy in Switzerland in 2000, Kahn and Matsusaka (1997) and Burkhardt and Chan (2017), who estimate demand for environmental policy (but not carbon pricing) using votes on a wide range of initiatives in California, and Holian and Kahn (2015), who study a failed 2010 vote to repeal California's cap-and-trade program (but not a carbon tax).⁴

In addition to our unique focus on a US carbon tax, we make several methodological contributions relative to these papers. First, we measure ideology separately from partisanship using votes on a dozen wide-ranging social, economic, and environmental policy issues. Our approach is related to that of Bornstein and Lanz (2008), who measure ideology using votes on five transportation-related referendums. We show that our simple measure of left-right ideology is highly predictive of voting on the carbon tax. Second, we go beyond merely explaining vote shares by measuring the overall willingness to pay (WTP) for the policy in dollars. Only one other study does this: Burkhardt and Chan (2017), who estimate median WTP for environmental policy (but not carbon pricing) in California as a whole. We build on this study by estimating WTP for a carbon tax for every precinct in our sample and showing the full distribution across precincts. This enables us to discover vast heterogeneity in WTP across precincts. In addition, we decompose WTP into four components, which allows us to show directly that ideology explains most of the variation in WTP across precincts, while energy tax incidence explains very little. This heterogeneity matters both for understanding distributional impacts and for campaign strategy, for example, fundraising and targeting marginal voters. Third, we show based on a bounding exercise that tax incidence also explains relatively little of the variation in WTP across individuals. Thus, we

4. Burkhardt and Chan (2017) also estimate preferences for nonenvironmental public goods, such as children's hospitals, while Holian and Kahn (2015) study a second low-carbon policy in California: high-speed rail.

are able to say more about individual-level heterogeneity than these other studies, which also rely on aggregate data. Finally, we are the first to use out-of-sample forecasts to illustrate the weak prospects for a carbon tax in other states.

We also contribute to a broader social science literature that studies the determinants of public opinion surrounding climate change, including preferences for carbon regulation and carbon taxes in particular. Egan and Mullin (2017) thoroughly review the survey literature, showing that partisan affiliation (Democrat vs. Republican) is the single most important driver of support for climate regulation and that the partisan gap has only widened in recent decades.⁵ Our main contribution here is to use actual voting data on a carbon tax. Like the survey literature, we find that partisanship is an important driver of support for a carbon tax.⁶ However, we show that it is not partisanship per se that mainly drives support but rather political ideology—which is strongly correlated with partisanship but measures something distinct. Further, our results offer a cautionary tale for the ability of surveys to predict the political viability of carbon taxes: carbon taxes failed at the polls even though they were popular in surveys, with 71% of registered voters in a 2018 Yale survey saying they support making fossil-fuel companies pay a carbon tax.⁷ Of course, researchers can frame their survey questions however they want, and respondents answer hypothetically in a low-information environment, free of competitive political messaging. In this setting, respondents might not know that a tax on oil companies will lead to higher gasoline prices or what trusted leaders would recommend. But in an actual election, ballot language is constrained, people are voting for real, and voters are bombarded with all kinds of information, including objective analysis and motivated campaign messaging that emphasizes both positives and negatives.

Finally, we contribute to a literature that studies the political feasibility and durability of carbon regulation, including the role of revenue recycling. See Rabe (2018) for a thorough review. In economics, two papers use postelection surveys to analyze failed Swiss referendums to tax nonrenewable energy. Thalmann (2004) finds that voters in 2000 prefer spending on green projects to cuts in other taxes, but the energy tax rates also differed across these policies (lower for green projects), muddying the

5. Egan and Mullin (2017) further emphasize (i) the importance of partisanship as a moderating variable for the effects of education and framing, (ii) that beliefs are not particularly susceptible to new information, and (iii) that belief in climate change does not necessarily lead to support for policy action.

6. This result is consistent with Holian and Kahn (2015), who find partisanship to be an important driver of actual voting on California's cap-and-trade policy in 2006. In contrast, Kahn and Matsusaka (1997) find that partisanship is not an important driver of voting on California's other forms of environmental regulation in the 1970s, 1980s, and 1990s.

7. See <https://climatecommunication.yale.edu/wp-content/uploads/2018/05/Global-Warming-Policy-Politics-March-2018.pdf>.

comparison of their revenue-recycling schemes.⁸ Meanwhile, Carattini et al. (2017) find that voters in 2015 prefer green projects but switch to preferring hypothetical lump-sum rebates when informed in a choice experiment about the progressive distributional effects. In Canada, Mildener et al. (2022) find that survey respondents in 2019 systematically underestimate the size of their federal carbon tax rebates. Conservatives underestimate rebates more than liberals and show no increase in policy support when properly informed but, rather, decrease in their belief that they receive more than they pay. In France, Douenne and Fabre (2022) find that survey respondents in 2019 would reject a carbon tax and dividend, in part because they are misinformed about the distributional and environmental effects. In the United States, survey respondents also prefer green projects to lump-sum transfers or reductions in other taxes (Amdur et al. 2014; Kotchen et al. 2017).

Our study is unique in using actual voting data for two policies that differ mainly in how they propose to use carbon tax revenue. We compare two different flavors of a carbon tax in back-to-back election cycles in the same state. Thus, we provide the most direct evidence to date on how voters perceive and react to different revenue-recycling schemes in a real-world electoral environment—with the obvious limitation that we observe only two votes and that other factors could have been at play. Our WTP framework implies that voters on average are not convinced that a revenue-neutral carbon tax will compensate them for higher energy prices—and do not much value the spending in the green-spending version of the carbon tax, if at all. Yet, beyond the average voter, we find that liberal voters prefer green spending, while conservative voters prefer the revenue-neutral policy, implying that revenue recycling is more of an ideological than a pocketbook issue in the eyes of voters. Thus, a carbon tax initiative with green spending will tend to do better in a liberal state like Washington—but still not well enough to get the initiative passed.

The rest of this paper proceeds as follows. Section 1 details the life of I-732 from its conception to the November 2016 election—and then the life of follow-on initiative I-1631 through its failure in November 2018. Section 2 describes our data sources. Section 3 explores the relationship between ideology and support for carbon taxes using precinct-level voting data. Section 4 tests whether economic incidence has a detectable impact on support for a carbon tax and estimates WTP for the policy's attributes. Section 5 forecasts vote shares for the two carbon taxes in other states based on our precinct-level regressions from within Washington. Finally, section 6 concludes with a discussion of lessons-learned and avenues for further research.

1. WASHINGTON STATE'S TWO CARBON TAX PROPOSALS

The I-732 campaign was spearheaded by Carbon Washington—a small grassroots organization led by Yoram Bauman, a professional stand-up comedian and PhD

8. Bornstein and Lanz (2008) study the 2000 referendums using actual voting data by municipality but do not focus on the political trade-offs due to differences in revenue recycling.

economist by training.⁹ Carbon Washington's strategy was to appeal to political moderates and conservatives, as well as liberals, through a revenue-neutral policy and targeted redistribution of carbon tax revenue. Thus, in addition to imposing a carbon tax, their policy would have reduced the state sales tax (to benefit all voters), expanded the EITC (to address distributional concerns), and reduced taxes on manufacturing businesses (to mitigate opposition from energy-intensive industry). Carbon Washington patterned this revenue-neutral "carbon tax swap" after a similar policy adopted by British Columbia in 2008. Meanwhile, the state's big player in carbon regulation over many years was the Washington Alliance for Jobs and Clean Energy (henceforth "Alliance"). The Alliance comprised a broad range of environmental, labor, and social justice advocacy groups. These members included important and well-resourced national-level environmental groups, such as Sierra Club and National Resources Defense Council, along with various state and local environmental groups. The labor and social justice groups reflected a similar range of national, state, and local advocacy groups—again including many heavy hitters such as the AFL-CIO.¹⁰ The Alliance's strategy was to explicitly tie carbon regulation to a program of spending on green jobs, improved health, and climate adaptation in low-income, historically disadvantaged communities. Thus, from the Alliance's perspective, any tax revenue should be targeted directly to these priorities.

The divergence between Carbon Washington and the Alliance highlights—in microcosm—a strategic fork in the road for would-be crafters of climate policy: appeal to moderates and conservatives through a revenue-neutral, market-based policy, or double-down on the left by spending revenue on the issues and identity groups that liberals care about. Carbon Washington turned right—or rather, aimed for the middle—while the Alliance veered left.

Carbon Washington unexpectedly gathered over 350,000 signatures, more than enough to put I-732 on the ballot. The Alliance approached Carbon Washington to negotiate a policy compromise that would satisfy both groups and that would, in the Alliance's view, do better at the polls. However, after a complicated discussion and some miscommunication, the two groups were unable to reach a compromise. Carbon Washington proceeded to the polls with I-732, while most members of the Alliance either actively opposed or—like the Sierra Club—declined to support I-732 (see app. A for details on who supported and opposed the two initiatives; apps. A–I are available online). The reasons for this opposition varied across groups but essentially boiled down to four issues: (1) concern that I-732 might lose revenue and put other

9. Bauman styles himself as "The world's first and only stand-up economist." See <http://standupeconomist.com>.

10. See the Alliance's web page for a statement of principals and list of members: <https://jobs-clean-energy-wa.com>.

programs at risk;¹¹ (2) a belief that revenue should be spent on issues important to the coalition, such as green jobs and climate adaptation; (3) a belief that such spending schemes polled better; and (4) the perception that Bauman and Carbon Washington failed to engage the broader social and environmental justice community in the design of I-732. After I-732 failed in 2016, the Alliance followed through in crafting and campaigning for I-1631 two years later.

Table 1 summarizes the key provisions of I-732 and I-1631. On the tax side, the two policies are similar: a carbon tax starting at \$15/tCO₂ and then rising gradually. However, on the revenue-recycling side, the two policies differ sharply. I-732 aims to be revenue-neutral, devoting most tax revenue to a reduction in the state sales tax and an expansion of the EITC—to mitigate impacts on low-income households.¹² In contrast, I-1631 allocates 95% of tax revenue to green projects and 5% to local communities. I-1631 appeals more to liberals, since it spends revenue on green investment, while I-732 appeals more to conservatives or moderates, since it is revenue neutral. This is because a major element of disagreement between the left and the right is the size and scope of government. Of course, both policies are broadly attractive to liberals, who prioritize immediate action on climate change and tolerate taxes more than conservatives.

The Alliance hoped that I-1631 would outperform I-732 for at least two reasons. First and most importantly, they thought that progressives would support using revenue for clean energy, while a revenue-neutral measure would alienate those progressives and win very few conservatives. Second, the I-1631 ballot language avoids the dreaded word “tax” and instead describes a “fee” on carbon.

In the following sections, we show that the Alliance’s forecast was at least partly correct: I-1631 performed substantially better than I-732 in liberal precincts—and somewhat worse in conservative ones. These outcomes are consistent with objective differences in revenue recycling. However, we cannot rule out other contributing factors. For example, forward-looking liberal voters in 2016 might have hoped for a better policy in the near future, which may have dampened their enthusiasm for I-732.¹³ The reported frictions between Carbon Washington and the Alliance may have reinforced these feelings. In addition, support for carbon taxes has been trending upward over time,

11. Independent estimates of the revenue impacts varied, highlighting this uncertainty: the Washington Office of Financial Management projected a 0.95% decrease in state revenue; Carbon Washington projected a 1.1% to 1.6% increase; and the Sightline Institute projected a –0.37% decrease. See Ballotpedia: [https://ballotpedia.org/Washington_Carbon_Emission_Tax_and_Sales_Tax_Reduction_Initiative_732_\(2016\)](https://ballotpedia.org/Washington_Carbon_Emission_Tax_and_Sales_Tax_Reduction_Initiative_732_(2016)).

12. Some may argue that the EITC increase is not a revenue-neutral tax cut but rather an increase in spending.

13. This mechanism implies that the difference between the support for I-732 and I-1631 is larger than what would be predicted by differences in revenue recycling alone. To the extent that voters listen to elite cues instead of only paying attention to policy design, it could also be that liberal elites were able to depress support for I-732 and boost support for I-1631.

according to the Yale Climate Opinion survey cited in the introduction, and this trend may be stronger among liberals. The surge of enthusiasm among Democratic voters (“blue wave”) in 2018 may have further contributed to the divergent outcomes. Finally, total campaign spending in 2018 was \$48 million—more than 10 times larger than in 2016—and dominated by the “no” campaign, which outspent the “yes” campaign by a factor of two. Meanwhile, total spending in 2016 was just \$4.6 million and dominated by the “yes” campaign, which outspent the “no” campaign by a factor of two (see table 1). These differences in spending may have influenced both the overall levels and patterns of support for the 2016 and 2018 policies.

2. VOTING DATA

In this section, we describe our data sources and procedures. Our main data come from the State of Washington Secretary of State (WA SOS) and record precinct-level election results from Washington State in the November 2016 and 2018 general elections.¹⁴ These data record the total number of votes cast for various candidates to elected office as well as total votes cast for and against various statewide ballot measures. We use these data to calculate—for each precinct—the share voting “yes” (vs. “no”) on the two carbon taxes (I-732 in 2016 and I-1631 in 2018) as well as the two-party share voting Republican in the 2016 US presidential election. For robustness checks, we also calculate Republican share separately for 2016 and 2018 based on the two US Senate elections in those years. We use Republican vote share to measure partisanship. In addition, the WA SOS data separately record the total number of registered voters and ballots cast in 2016 (these data are not available for 2018). We use these data to calculate the share of registered voters who cast ballots in 2016 (turnout) as well as the share of registered voters who recorded a vote either for or against the carbon tax in 2016 and 2018 (carbon tax turnout). The latter measure of turnout is generally lower due to roll-off (when voters turn in a ballot without making all possible choices). For comparison, we also calculate similar turnout measures for other initiatives and elections.

To measure latent ideology, we conduct a principal component decomposition of the precinct-level vote shares on 12 other statewide ballot measures from November 2016 and 2018. These measures cover a wide range of social, economic, and procedural issues.¹⁵ We find that the first component in this decomposition explains

14. Available at <https://www.sos.wa.gov/elections/research/election-results-and-voters-pamphlets.aspx>.

15. See app. B for details on each measure, along with scatter diagrams showing precinct-level vote shares on each measure versus the presidential vote. Turnout on these measures (and the carbon tax) is uniformly high, is strongly correlated across precincts, and follows the same pattern relative to ideology: uniformly high for liberal and conservative precincts with a slight dip among moderates. These results are consistent with the same voters (or types of voters) voting on all of the ballot measures, which supports our issues-based approach to measuring a precinct’s latent ideology. See app. D for summary statistics and descriptive figures on turnout.




Figure 1. Precinct-level conservative ideology versus Republican presidential vote. This figure plots our constructed index of conservative ideology versus the Republican party vote share in the 2016 presidential election for 6,219 precincts in Washington State. Our index of conservative ideology is based on votes for ballot measures other than the carbon tax in 2016 and 2018 (see app. B). Data source: Washington Secretary of State.

80% of the variation in voting on these ballot measures.¹⁶ We interpret this component as an index measure of issues-based latent ideology. We linearly transform this index to range from zero (most liberal) to one (most conservative). In our results, we label our measure of ideology as “conservatism” to facilitate interpretation, since ideology in principle could go from left to right or vice versa. Figure 1 plots our resulting measure of conservative ideology versus the share voting Republican in each precinct (our measure of partisanship). The figure shows that, while these two measures are highly correlated (note the correlation coefficient of 0.97 in table 2), they are not perfectly correlated. Indeed, at each level of Republican vote share, there are both relatively conservative and relatively liberal precincts. Thus, we are able to estimate support for the carbon tax conditional on both measures below.

We match these voting data to US Census data as follows. First, we obtain data from WA SOS that provide the distribution of each voting precinct’s total population in 2010 across Washington’s more than 100,000 census blocks (based on block

16. The eigenvalue on the first component is 9.60, which implies that this component explains $9.60/12 = 80\%$ of the variation spanned by the 12 ballot measures.

centroids).¹⁷ We use these data to calculate, for each precinct, the share of the population living in each of Washington's roughly 4,800 census block groups. Second, we match these population shares to US Census block group aggregate data. Third, for each precinct, we calculate the population-weighted averages of the block-group-level data. Finally, we match these precinct-level weighted averages of the underlying block-group-level data to precinct-level election data.¹⁸

US Census data come from American Community Survey (ACS) five-year estimates for 2012–16. In constructing population-weighted averages of block-group-level data, we do not rely on block-group-level medians. Rather, we rely on block-group-level shares of people, households, or workers that fall into narrow categories. For example, we use the share of people age 40–44 or the share of households with incomes of \$50,000–\$59,999 rather than median age and median income. This approach leads to more sensible data aggregation and allows us to more flexibly model the relationship between census covariates and support for a carbon tax.¹⁹ Coefficients on these variables (e.g., population share age 40–44) intuitively have the same interpretation as those on a dummy variable (e.g., probability shift for a voter age 40–44). Our detailed census variables measure population shares for car commute time (for people age 16+); education (for people age 25+); industry (for workers age 16+); vehicle ownership, annual income, lives in an urbanized area, owner occupied, home value, and home size (for households); and age, race, and gender (for all people). We also control for mean household size (total population divided by number of households). See appendix C for our full set of census variables.

Given the monocentric city model (von Thünen 1826; Fujita 1989), we expect suburban and rural households to occupy larger single-family homes and to commute longer distances via car, such that they will consume more energy than urban households and be more affected by a carbon tax. Our controls for home size, owner occupied, vehicle ownership, and car commute time capture these important differences between

17. We thank Nicholas Pharris at WA SOS for providing these data.

18. We calculate for 2016 that 34% of precincts overlap with one block group, 37% overlap with two, 18% overlap with three, 7% overlap with four, 3% overlap with five, and the remaining 1% overlap with six, seven, or eight block groups. In total, 89% of precincts overlap with three or fewer block groups.

19. Consider a precinct that overlaps with multiple block groups indexed by $j = 1, 2, \dots, J$. Let μ_j be the share of the precinct's population in block group j with $\sum_{j=1}^J \mu_j = 1$. Let θ_j be the share of block-group j 's population in some demographic category (e.g., age 40–44). Finally, assume that a block group's population is distributed homogeneously through its geographic area (e.g., the age 40–44 population is not concentrated in one part of the block group or another). Then the share of the precinct in the given demographic category is $\sum_{j=1}^J \mu_j \theta_j$, i.e., the population-weighted average of the block-group shares. In contrast, the median value of some demographic variable (e.g., median age) in a precinct is not in general the population weighted-average of the block-group-level medians.

urban, suburban, and rural areas. We also control directly for the share of households in the surrounding county that live in an urbanized area. In some models, we also control for county fixed effects, which absorb this variable and any other observed and unobserved differences across counties, including housing and population density.

We make several sample restrictions. First, we limit our analysis to precincts that did not experience boundary changes between 2016 and 2018.²⁰ Our focus on these precincts allows us to directly measure changes in support for the carbon tax between 2016 and 2018, and to measure a precinct's ideology using ballot initiatives from both years. This restriction omits 9% of precincts. Second, we limit our analysis to precincts with at least 50 votes cast in the November 2016 presidential election to more precisely estimate latent partisanship and ideology.²¹ This restriction omits an additional 4% of precincts. Finally, we focus on observations with complete census data. This restriction omits just 0.06% of precincts. In the end, we are left with a sample size of 6,219 precincts for 12,438 pooled observations across our two elections (2016 and 2018). Our final sample includes 87% of the 2016 precincts representing 88% of the state's population and 89% of the voter turnout.

Table 2 presents summary statistics and pairwise correlation coefficients for the election-based variables measuring support for the carbon tax in 2016 and 2018, along with Republican share and conservative ideology. We weight by the total presidential vote so that our sample statistics more closely reflect the underlying population of voters rather than precincts. Note that support for the carbon tax is strongly negatively correlated with both Republican share and conservative ideology (correlation coefficients ranging from -0.93 to -0.98). See appendix C for summary statistics on the full set of census demographic variables, including important correlates of household carbon footprint, such as vehicle ownership and home size.

3. WHO SUPPORTS A CARBON TAX?

In this section, we investigate the correlates of support for the carbon tax using aggregate precinct-level voting data for I-732 and I-1631. We focus on how ideology and

20. Precincts experience boundary changes when they split (creating a new precinct and corresponding precinct code), merge (eliminating an existing code), or shift boundaries to rebalance population (so that old codes correspond to different geographic areas). WA SOS maintains GIS shapefiles that record precinct boundaries for each year. At our request, Nicholas Pharris at WA SOS used these shapefiles to match each of the state's more than 100,000 census blocks to precinct boundaries in 2016 and 2018 (based on block centroids). Thus, we are able to identify precincts that experience boundary changes based on census blocks that match to different precinct codes in 2016 and 2018.

21. Small precincts yield noisy estimates of a precinct's latent partisanship and ideology, leading to multiple cases of 0% and 100% vote shares. These extreme cases skew our regression results. We found that omitting predicts with fewer than 50 presidential votes cast eliminated

Table 2. Summary Statistics and Correlation Coefficients for Election-Based Variables

	Mean	SD	732	1631	Pooled	Rep.	Cons.
Yes on I-732 (2016)	.409	.125	1.00				
Yes on I-1631 (2018)	.434	.172	.94	1.00			
Yes on carbon tax (pooled)	.421	.151	.95	.96	1.00		
Republican	.412	.192	-.94	-.95	-.93	1.00	
Conservatism	.530	.177	-.95	-.98	-.95	.97	1.00

Source. Data from the Washington Secretary of State.

Note. The left panel of this table reports means and standard deviations for our estimation sample of 6,219 precincts (12,438 observations in our pooled sample), weighted by the total votes cast for the Republican or Democratic candidate in the 2016 presidential election. The first two variables measure share voting “yes” on the carbon taxes in 2016 and 2018. The third variable stacks these outcomes in a pooled sample of precinct-year observations. Republican measures the share voting for the Republican (vs. Democrat) in the 2016 presidential election. Conservatism measures ideology on a 0–1 index. The right panel reports pairwise correlation coefficients for all five variables, again weighted by total votes cast for the Republican or Democratic candidate for president in 2016.

partisanship predict votes for the carbon tax at the precinct level, and we document the policy preferences that ideology captures. We then compare the 2016 and 2018 carbon tax initiatives in terms of their appeal to different ideological segments of the voting population. Finally, we show that our results are robust to several data issues and modeling choices.

3.1. Explaining Overall Support Using Pooled 2016 and 2018 Data

Figure 2 shows the share voting “yes” on each of the two carbon taxes by decile of the Republican vote share in the 2016 presidential election.²² Deciles for 2018 (solid line) are constructed such that they each capture the same number of votes cast for or against the carbon tax in 2018, rather than the same number of precincts—and similarly for 2016 (dashed line). Thus, the overall share voting “yes” on the carbon tax in a given year can be read visually as the average height of the corresponding line. This figure illustrates clearly that support for the carbon tax falls as the Republican share increases. The relationship is stronger in 2018 (solid line) than in 2016 (dashed line). Indeed, the most liberal precincts (on the left) tend to prefer the 2018 policy, while the most conservative precincts (on the right) tend to prefer the 2016 policy.²³ In principle, the negative correlation between support for the carbon tax and Republican share in

all such extreme cases. We omit these precincts prior to calculating and rescaling the latent ideology measures described above.

22. See app. E for scatter diagrams showing the underlying precinct-level data.

23. We repeat this graphical analysis using conservative ideology in place of the Republican vote share and arrive at similar conclusions. See app. E.




Figure 2. Vote share on carbon taxes in 2016 and 2018 as a function of Republican presidential vote. This figure plots the “yes” shares on I-1631 in 2018 (solid line) and I-732 in 2016 (dashed line) by decile of the Republican party vote share in the 2016 presidential election. Each decile contains precincts that together add up to one-tenth of votes cast in 2016 and 2018, respectively. Deciles for 2018 are constructed from precinct-level data in the following way: sort precincts from the lowest to the highest Republican vote share and then determine decile cut-offs. Deciles for 2016 (dashed line) are constructed similarly. Thus, the overall vote share can be visualized as the average height of the points. Data source: Washington Secretary of State.

this figure could reflect political ideology—or it might simply reflect the fact that Republicans tend to live in the suburbs, where people own more cars, have longer commutes, and reside in bigger houses, and would therefore see larger increases in energy costs under a carbon tax. Therefore, we also explore the relationship between support for a carbon tax and political ideology in regressions that control for these and other variables.

Table 3 presents ordinary least squares (OLS) regression results for precinct-level carbon tax vote shares conditional on ideology, partisanship, and demographics. We continue to weight by the total presidential vote in each precinct; unweighted results are nearly identical. These regressions pool voting outcomes from 2016 and 2018; each regression includes a 2018 dummy, but the other explanatory variables are purely cross-sectional.²⁴ Conservative ideology is a highly significant predictor of support for the

24. This data structure explains why the coefficients and standard errors on the 2018 dummy are the same in each column.

Table 3. Predicting the Carbon Tax Vote Share at the Precinct Level (pooled 2016 and 2018)

	Ideology (1)	Party (2)	+Census (3)	+Ideology (4)	+County FEs (5)	+Initiatives (6)
Conservatism	-.814*** (.013)			-.669*** (.019)	-.665*** (.018)	
Republican		-.730*** (.026)	-.635*** (.021)	-.124*** (.018)	-.139*** (.018)	-.046** (.016)
2018 Vote	.026 (.014)	.026 (.014)	.026 (.014)	.026 (.014)	.026 (.014)	.026 (.014)
Observations	12,438	12,438	12,438	12,438	12,438	12,438
R ²	.913	.870	.906	.926	.929	.932

Source. Data from Washington Secretary of State and US Census.

Note. This table presents coefficient estimates from pooled precinct-level OLS regressions modeling the share voting "yes" for the carbon tax in 2016 and 2018 as a function of ideology, demographics, and other factors. Conservatism measures ideology on a 0–1 index, computed from the precinct's vote shares on 12 ballot measures in 2016 and 2018 (see app. B). Republican measures the share voting for the Republican (vs. Democratic) party candidate in the 2016 US presidential election. 2018 Vote is an indicator for the 2018 carbon tax (I-1631). Models 3–6 build cumulatively on model 2. Model 3 controls for detailed census variables (no. of vehicles, commute time by car, industry, home value, no. of rooms, income, gender, age, race, education, household size, urban share, and owner-occupied share). Model 4 then adds ideology. Model 5 then adds county fixed effects. Finally, model 6 replaces ideology with vote shares for the 12 individual ballot initiatives. For all models, precincts are weighted by the total votes cast for the Republican or Democratic candidate in the 2016 presidential election. Standard errors are clustered by county (39 clusters).

* Significant at the 5% level.

** Significant at the 1% level.

*** Significant at the .1% level.

carbon tax with an *R*-squared of 91.3% (col. 1).²⁵ On average, support for the carbon tax falls by 81 percentage points, moving from the most liberal to the most conservative precinct. Meanwhile, the share voting Republican in the 2016 presidential election is also a highly significant predictor of support for the carbon tax, but its overall predictive power is slightly lower with an *R*-squared of 87% (col. 2). On average, support falls by 73 percentage points, moving from a precinct that votes 100% Democratic to a precinct that votes 100% Republican. Adding a slew of census demographics to the Republican vote share bumps the *R*-squared up to 90.6% (col. 3), but conservative ideology alone still has slightly higher predictive power (col. 1). Adding conservative ideology on top of the Republican vote share and census demographics further boosts the *R*-squared to 92.6% (col. 4). More interesting, however, is the fact that the coefficient on Republican share shrinks by a factor of five, while the coefficient on conservative ideology remains quite high. These variables both range from 0 to 1 and have nearly

25. By itself, the 2018 dummy has an *R*-squared of just 0.7%. Thus, as an approximation, we wholly attribute the *R*-squared values in this table to the other variables.

identical variance (see fig. 1 and table 2), facilitating a direct comparison of their regression coefficients. Thus, if we were to treat this as a causal model (i.e., no correlated omitted variables), the results would imply that ideology is the dominant driver of support for a carbon tax. In principle, there could be some other variable—highly correlated with ideology—that is the true, underlying driver of support for a carbon tax, and controlling for this hypothetical variable would substantially undermine the apparent contribution of ideology. It is difficult to know what such a variable or set of variables might be, given that we already control for partisanship and an exhaustive set of detailed demographic controls. Yet because we cannot rule out this possibility, our results are merely consistent with the idea that ideology is the main driver of policy preferences and votes. Either way, the predictive value of partisanship is substantially weakened when we control for ideology, while the reverse is not true.²⁶ Likewise, demographic variables (including key proxies for energy tax incidence, such as vehicle ownership) contribute relatively little predictive power to a model that already controls for ideology and partisanship.

As suggested above, one potential concern is that we have omitted important variables that are correlated with partisanship and ideology, undermining our interpretation of these variables as the dominant drivers of support for a carbon tax. We are mainly concerned about omitted correlates of energy use, such as vehicle fuel economy or local climate conditions. To address this concern, we include county-level fixed effects, which would help control for any such variables that differ across counties. We find that the coefficients on ideology and partisanship do not change (col. 5). Overall, these results confirm that the strong correlation we observe in figure 2 between support for a carbon tax and political ideology is not primarily driven by differential energy tax incidence for liberals as compared to conservatives.

Recall that we measure ideology by exploiting the presence of 12 statewide ballot measures considered by Washington voters in 2016 and 2018 (see app. B). We use vote shares on these measures across precincts to calculate a precinct-level index of political ideology based on the first component from a principal component decomposition. Creating an index allows us to capture the effects of ideology in a single, intuitive measure. This approach is based on the assumption that voting on these other measures is also ideological in nature and consistent across issues.

To probe deeper, we estimate an additional regression model (6) in table 3. This model is identical to model 5 in the table but replaces the single ideological index (Conservatism) with 12 variables measuring the “yes” vote shares on each of the underlying

26. This pattern shows that we can separately identify the predictive effect of ideology vs. partisanship, despite the high correlation between the two. If ideology and partisanship were essentially the same, we would expect to see highly imprecisely estimated coefficients when including both.




Figure 3. How votes on other ballot measures predict the vote on the carbon tax in 2016 and 2018. This figure plots coefficients from the regression reported in column 6 of table 3. This is a pooled OLS regression modeling the share voting “yes” for the carbon tax in 2016 and 2018 as a function of the share voting “yes” on other ballot measures. The measures on regulate deadly force (I-940), gun control (purchase) (I-1639), no local grocery taxes (I-1634), and repeal pipeline oil tax (Advisory vote 19) were from 2018. The measures on gun control (ownership) (I-1491), ID theft penalty (I-1501), minimum wage (I-1433), rights for people only (I-735), campaign finance system (I-1464), earlier redistricting (Senate Joint Resolution no. 8210), repeal clean car subsidy (Advisory vote 15), and repeal dental tax (Advisory vote 14) were from 2016 (see app. B for details). Point estimates are represented by dots, while 95% confidence intervals (based on standard errors clustered by county) are represented by horizontal lines. Control variables include the share voting Republican (vs. Democrat) in the 2016 US presidential election, the 2018 vote dummy, detailed census variables (no. of vehicles, commute time by car, industry, home value, no. of rooms, income, gender, age, race, education, household size, urban share, and owner-occupied share), and county fixed effects. Data source: Washington Secretary of State and US Census.

ballot measures. Figure 3 reports the coefficient estimates on these variables. Support for repealing a clean car subsidy and repealing an oil pipeline tax both correlate negatively with support for a carbon tax. This is perhaps not surprising, given that these policies also have direct environmental implications. What is noteworthy for our model of ideology, however, is that support for tighter gun control and higher minimum wages—largely or entirely unrelated social and economic policies—also correlate (positively) with support for a carbon tax. This result is consistent, however, with the notion

of ideology acting as a psychological constraint on beliefs and policy preferences (Converse 1964).

3.2. Comparing Two Flavors of the Carbon Tax in 2018 versus 2016

Above, we study the overall level of support for carbon taxes. Here, we explore differences in support for a green-spending policy (I-1631 in 2018) relative to a revenue-neutral policy (I-732 in 2016). To begin, figure 4 plots the change in support for I-1631 relative to I-732 by decile of Republican vote share.²⁷ The figure shows that liberal precincts increased their support for I-1631 relative to I-732. At the same time, conservative precincts decreased their support. Specifically, I-1631 gained 3.29 percentage points in cumulative vote share relative to I-732 among more liberal precincts (deciles 1–6) and lost 0.65 percentage points among more conservative precincts (deciles 7–10) for a net gain of 2.64 percentage points.²⁸ We are hesitant to make any strong causal claims based on a before-after comparison of voting outcomes, since liberal versus conservative support could have shifted for other reasons. However, the pattern of these gains and losses is consistent with the hypothesis that both the progressive (green-spending) design of I-1631 as well as the moderate- and conservative-appealing design of I-732 worked as intended by their respective promoters—at least to a degree.

Next, we regress changes in the precinct-level vote share on ideology, partisanship, demographics, and other variables. Table 4 presents the OLS regression results—with positive coefficients indicating that a given variable is associated with an increase in vote share for the green-spending package (I-1631 in 2018) relative to the revenue-neutral policy (I-732 in 2016). Conservative ideology is a highly significant predictor of the change in vote share, with an *R*-squared of 51.4% (col. 1). On average, the green-spending policy loses 27.8 percentage points in relative support moving from the most liberal to the most conservative precinct. Meanwhile, the share voting Republican in the 2016 presidential election is also a highly significant predictor of the change in vote share, but its overall predictive power is lower with an *R*-squared of just 44.1% (col. 2). On average, the green-spending policy loses 23.7 percentage points in relative support moving from a precinct that votes 100% Democratic to a precinct that votes 100% Republican. Adding census demographics increases the *R*-squared but has virtually no effect on the coefficient for Republican share (col. 3). As above, the coefficient on Republican share shrinks dramatically when we control for conservative ideology, while the coefficient on ideology remains high (col. 4). Again, these inferences are robust to including county-level fixed effects (col. 5). Replacing the ideology variable with support

27. See app. E for a scatter diagram showing the underlying precinct-level data.

28. We repeat this graphical analysis using conservative ideology in place of the Republican vote share and arrive at similar conclusions. See app. E.




Figure 4. Change in vote share on carbon tax (2018 minus 2016) versus presidential vote (by decile). This figure plots changes in average “yes” shares (I-1631 in 2018 relative to I-732 in 2016) by decile of the US presidential vote share (Republican) in 2016. Each decile contains precincts that together add up to one-tenth of votes cast in 2016 and 2018, respectively. Deciles for 2018 are constructed from precinct-level data in the following way: sort precincts from the lowest to the highest Republican vote share and then determine decile cutoffs. Deciles for 2016 are constructed similarly. Thus, the overall difference in vote shares between 2018 and 2016 can be visualized as the average height of the points. Data source: Washington Secretary of State.

for each of the ballot initiatives does not substantially change the coefficient on Republican (col. 6). Our regression results are thus consistent with figure 4, showing that more liberal districts prefer the green-spending version of a carbon tax, while conservatives prefer the revenue-neutral policy. Further, our results are consistent with the idea that ideology—interacting with differences in the perceived attributes of the two policies—is the main driver of the shift in support for I-1631 relative to I-732. In principle, there could be some other cross-sectional variable—highly correlated with ideology—that explains why liberals preferred the 2018 policy, such that controlling for this variable would render ideology irrelevant. But this does not seem especially plausible, given our controls for partisanship and demographics. Meanwhile, liberal versus conservative support for carbon taxes could have been shifting over time for reasons having little to do with the details of these two policies, such as changes in the amounts spent on campaigns. Regardless, we can say the predictive value of partisanship is severely undermined when we control for ideology, while demographics contribute relatively little predictive value.

Table 4. Predicting the Change in the Carbon Tax Vote Share (2018 minus 2016)

	Ideology (1)	Party (2)	+Census (3)	+Ideology (4)	+County FEs (5)	+Initiatives (6)
Conservatism	-.278*** (.011)			-.268*** (.047)	-.249*** (.019)	
Republican		-.237*** (.013)	-.224*** (.015)	-.020 (.044)	-.049* (.023)	-.033 (.024)
Observations	6,219	6,219	6,219	6,219	6,219	6,219
R ²	.514	.441	.618	.634	.689	.755

Source. Data from Washington Secretary of State and US Census.

Note. This table presents coefficient estimates from OLS regressions modeling the change in share voting “yes” for the carbon tax in 2018 vs. 2016 as a function of ideology, demographics, and other factors. Conservatism measures ideology on a 0–1 index. Republican measures the share voting for the Republican (vs. Democratic) party candidate in the 2016 US presidential election. Models 3–6 build cumulatively on model 2. Model 3 controls for detailed census variables (no. of vehicles, commute time by car, industry, home value, no. of rooms, income, gender, age, race, education, household size, urban share, and owner-occupied share). Model 4 then adds ideology. Model 5 then adds county fixed effects. Finally, model 6 replaces ideology with vote shares for the 12 individual ballot initiatives. For all models, precincts are weighted by the total votes cast for the Republican or Democratic candidate in the 2016 presidential election. Standard errors are clustered by county (39 clusters).

* Significant at the 5% level.

** Significant at the 1% level.

*** Significant at the .1% level.

3.3. Robustness Checks

This subsection discusses the sensitivity of our results to sample selection, shifts in partisanship over time, ecological regression bias, and potential bias caused by differences in voter turnout by party, ideology, and demographics.

3.3.1. Sample Selection

Recall from section 2 that our estimation sample omits approximately 9% of precincts due to changing precinct boundaries and an additional 4% of precincts with small vote counts. Regression on this selected sample will yield consistent estimates if the full-sample error terms (u) are uncorrelated with each of the covariates (x_j) in the selected sample: $E(su \cdot sx_j) = 0$, where s is a selection indicator equal to 1 if the observation is included and 0 otherwise (Wooldridge 2020). We test this condition directly by (i) regressing support for the carbon tax in 2016 on Republican share using the full cross-sectional sample of 2016 precincts (analogous to regression 2 in table 3), (ii) generating the residuals, and (iii) regressing these residuals on Republican share in the selected sample. We use Republican vote share for this analysis because Republican share is never missing, while ideology is missing for the 2016 precincts that later experienced boundary changes. The resulting coefficient is 0.0030 (equal to the difference in coefficients

between the full and selected samples in 2016) and statistically insignificant with a heteroskedasticity robust standard error of 0.0033. Meanwhile, the model R -squared is just 0.0002, and an F -test for the model's overall significance fails to reject that the model explains nothing (p -value of .3082). When we repeat this procedure adding our full set of census demographic controls and county fixed effects, coefficients on Republican share are small and insignificant, model R -squared values remain near zero, and F -tests for overall significance still fail to reject (p -values near 1). Thus, we conclude that the sample selection bias here is minuscule.

3.3.2. Shifts in Partisanship

We use a single, cross-sectional measure of partisanship based on the Republican presidential vote share in 2016 to explain support for the carbon tax in both 2016 and 2018. This approach could yield biased estimates if partisanship shifted over time within precincts, perhaps due to differential changes in turnout for Republicans versus Democrats in a presidential versus midterm election. To explore this issue, we measure partisanship using the US Senate races in 2016 and 2018, which allows us to match Republican share in 2016 to the carbon tax vote in 2016 (I-732) and Republican share in 2018 to the carbon tax vote in 2018 (I-1631). We replicate figures 2 and 4 based on this alternative, year-specific measure of partisanship. We find that the results look nearly identical (see app. E). Likewise, we replicate regression 2 from table 3 using this year-specific measure of Republican. We obtain a coefficient of -0.853 (see app. F for this and other regression results based on the Senate vote). To parallel our current approach, we then use the Republican share in 2016 only and obtain a nearly identical estimate of -0.885 . Meanwhile, when we include both the new year-specific measure of Republican and our old cross-sectional measure of Conservatism simultaneously, we continue to find that the coefficient on Conservatism is larger. This gap only magnifies when we use a cross-sectional measure of partisanship, whether the Republican share in 2016 or a weighted average of 2016 and 2018. Thus, our qualitative results are confirmed.

3.3.3. Ecological Bias

To interpret the coefficient on Republican share as the effect of being an individual Republican voter on support for the carbon tax, all else equal, we must assume that there is no ecological bias. Specifically, we must assume that Republican voters across precincts have the same opinion about the carbon tax (constancy assumption). If Republican voters in more Republican precincts dislike the carbon tax more, then the coefficient on Republican share—while still capturing the effect of partisanship in a generic sense—will overstate the effect of being an individual Republican voter drawn at random from the voting population. Ecological bias is less of a concern for the coefficient on ideology, because our measure of ideology captures aggregate policy preferences more finely, using votes on many other ballot measures. Individual voters differ continuously in their ideology. Thus, support for the carbon tax in a precinct will tend

to be higher if there are more liberals in the precinct—or if the liberals in the precinct are more liberal. Our aggregate measure of ideology implicitly captures both effects.

3.3.4. *Turnout*

We measure aggregate support for the carbon tax among people who actually vote. If turnout varies with partisanship, ideology, and demographics, then the regression results in table 3 may not reflect the underlying preferences of voters. For example, suppose we are interested in the preferences of young people. Young people are more likely to support carbon taxes (positive effect on percentage support) but less likely to vote (negative effect on percentage support). So will the correlation between percentage support and percentage young accurately reflect the preferences of young versus old? Or is this correlation biased? Lang and Pearson-Merkowitz (2021) show that the bias is less severe when turnout is high and the level of geographic aggregation is small, and our precinct-level data from Washington compare favorably to similar studies in both dimensions.²⁹ We explore this issue further in appendix G, where we develop a simple framework to assess whether differential turnout by party, ideology, and demographics is strong enough to severely bias our results. Our results suggest that the bias is likely small. But we cannot definitively rule out bias, given the simplifying assumptions of our framework and reliance on aggregate data for calibration. Microdata based on exit polling could potentially speak to these concerns but are no panacea. Exit polling itself is fraught with challenges. Representative sampling is hard, and nonresponse is a real problem. If respondents in exit polls are selected on unobservables, the polls will yield neither a representative sample of voters nor a representative sample of the underlying population, and the correlations based on exit polls could be biased. Meanwhile, our regressions based on aggregate data will at least remain valid as reduced-form correlations between voting outcomes and the underlying population.

29. Washington turnout in 2016 was 79% among registered voters and 61% among those eligible to vote. These values were 72% and 53% in 2018. See Washington turnout data: <https://www.sos.wa.gov/elections/voter-participation.aspx>. Meanwhile, three of the studies we cite use data from California (Kahn and Matsusaka 1997; Holian and Kahn 2015; Burkhardt and Chan 2017). These studies cover the years 1972–2010, during which time turnout averaged 62% among registered voters and 46% among those eligible to vote. These values were 60% and 44% in 2010, the year of California's cap-and-trade referendum, which is the focus of Holian and Kahn (2015). See California turnout data: <https://elections.cdn.sos.ca.gov/sov/2010-general/04-historical-voter-reg-participation.pdf>. These California studies use county (Kahn and Matsusaka 1997), zip code (Burkhardt and Chan 2017), and precinct-level (Holian and Kahn 2015) voting data. The other study uses municipality-level data from Switzerland in 2000, reporting voter turnout of just 45% (Bornstein and Lanz 2008).

4. TAX INCIDENCE AND WILLINGNESS TO PAY FOR A CARBON TAX

To what extent does personal tax incidence drive opposition to the carbon tax? Voters with the highest personal energy consumption—in particular, those consuming more gasoline to propel their cars or more electricity and natural gas to power and heat their homes—will tend to incur the largest direct costs from a carbon tax. Thus, we might expect lower support from such voters, all else equal.

To test this proposition, figure 5 reports the coefficient estimates on vehicle ownership (number of vehicles) and home size (number of rooms) from model 6 in table 3.³⁰ These variables proxy for personal energy consumption. Point estimates are represented by dots, while 95% confidence intervals are represented by vertical lines. Overall, our results support the idea that voters facing higher direct costs are less likely to support the carbon tax. The coefficients on the number of vehicles are all negative (top panel). Thus, having at least one vehicle is associated with lower support for the carbon tax. Surprisingly, having many vehicles is not associated with especially low support, even though microdata show that gasoline consumption tends to rise linearly with number of vehicles (see sec. H.2.1 of the appendix). However, note that fewer than 10% of households have four or more vehicles (see app. C). Meanwhile, the coefficients on the share of households with three or more rooms are all negative (bottom panel), while the coefficient on two rooms is almost exactly zero. Thus, relative to homes with just one or two rooms, larger homes are associated with lower support for the carbon tax. Moreover, support tends to decline as homes get larger, which is consistent with higher energy consumption for larger homes (see sec. H.2.2 of the appendix).³¹

What can this trade-off between tax incidence and support for the carbon tax tell us about voters' willingness to pay for these policies? To answer this question, we push our analysis several steps further by measuring the direct personal energy tax incidence—in dollars—that Washington's proposed carbon taxes would have had on Washington voters. We then estimate a structural discrete-choice random-utility model, relating the probability of voting yes on the carbon tax to dollars of energy tax incidence at the precinct level. Thus, we are able to convert vote shares into precinct-level mean WTP in dollars. This approach yields total WTP for the overall policy package, that

30. This model controls for the Republican vote share in the presidential election, the votes on the 12 individual ballot measures (other than the carbon tax), the 2018 dummy, the full suite of census demographic variables—including income, age, gender, and manufacturing industry, which may be necessary to control for the tax reductions and rebates embedded in the I-732 package—and county fixed effects.

31. Full results for the census demographic variables, available upon request, show that the coefficients on number of vehicles and number of rooms are not especially sensitive to the inclusion of county fixed effects nor to the use of all 12 individual ballot measures in place of the ideological index.




Figure 5. Predicting the carbon tax vote share by precinct in 2016 and 2018: coefficients on number of vehicles and number of rooms. These figures plot coefficients from the regression reported in column 6 of table 3. This is a pooled OLS regression modeling the share voting “yes” for the carbon tax in 2016 and 2018. The top panel plots coefficients on variables measuring the share of households with various numbers of vehicles available; the excluded category is households with zero vehicles. The bottom panel plots coefficients on variables measuring the share of households with various numbers of rooms in their home; the excluded category is homes with just one room. Point estimates are represented by dots, while 95% confidence intervals (based on standard errors clustered by county) are represented by vertical lines. Control variables include the share voting “yes” on 12 other ballot measures, the share voting Republican (vs. Democrat) in the 2016 presidential election, the 2018 vote dummy, detailed census variables (commute time by car, industry, home value, income, gender, age, race, education, household size, urban share, and owner-occupied share), and county fixed effects. Data source: Washington Secretary of State and US Census.




Figure 6. Decomposition of precinct-level mean willingness to pay for a carbon tax policy. EITC = Earned Income Tax Credit.

is, the perceived net benefits from the referendum passing (see the top level of fig. 6). This package includes the direct energy tax incidence due to higher gasoline and residential energy prices, which enters negatively in the net-benefits ledger, plus other policy impacts (see the middle level of fig. 6). These other policy impacts include the lower sales tax and expanded EITC (in 2016), local green-spending projects (in 2018), and environmental benefits accruing to Washington voters and others via reductions in local emissions and global greenhouse gases. Finally, we decompose WTP for these other policy impacts into components that can be explained by political ideology, observed demographics, and a residual (see the bottom level of fig. 6).

Our WTP analysis assumes, like all revealed preference studies, that the tax incidence we calculate for Washington voters mirrors their actual beliefs about the cost of the 2016 and 2018 policies. We model average support for a carbon tax by precinct. Thus, our assumption is that average beliefs in a precinct align with our calculations for that precinct's average tax incidence. As we discuss below, our calculations are likely to be quite accurate. Further, we think that voters would have been able to form a fairly accurate estimate of their own tax incidence, such that their average beliefs align reasonably well with our calculations, for three reasons. First, Washington has 100% mail-in ballots, which are sent to registered voters at least 18 days prior to election—along with voter guides that provide substantial detail on all statewide ballot initiatives. This gives voters time to learn about the various elections and contemplate their votes while gaining additional information. Second, prior research demonstrates that consumer beliefs and behavior are consistent on average with reasonable forecasts for energy prices (Anderson et al. 2013) and consumption (Anderson and Sallee 2016, 171–72). Third, as we describe below, web-based carbon tax calculators available at the time would have allowed Washington voters to calculate their personal energy tax incidence under both I-732 and I-1631 prior to voting, and our parameter assumptions are informed by these tax calculators. Of course, voters may not always have availed themselves of these resources and

may have used simple heuristics instead. Thus, we consider several possible sources of divergence between our calculations and voters' beliefs below.

4.1. Empirical Approach for Estimating Willingness to Pay

Our analysis proceeds in several steps. First, we calculate tax incidence related to household gasoline consumption and home energy consumption for the average voting-age adult (18+) in each precinct. We assume that voters perceive a \$25/tCO₂ carbon tax in 2016 and a \$15/tCO₂ carbon tax in 2018. These values are consistent with web-based tax burden calculators available to households at the time of the 2016 and 2018 votes as well as with contemporaneous reporting. Meanwhile, using household-level microdata from the 2017 National Household Transportation Survey (NHTS), we regress gasoline consumption on closely related household characteristics: number of vehicles, car commute time, number of household members, and an indicator for urbanization. Likewise, using microdata from the 2015 Residential Energy Consumption Survey (RECS), we regress home energy emissions on closely related characteristics: number of rooms, number of household members, and owner-occupied.³² We then apply the coefficients from these micro regressions to the corresponding precinct-level averages to estimate mean household-level carbon emissions for each precinct. Finally, we divide by the average number of voting-age adults per household in each precinct and multiply by the perceived carbon taxes in each year, to yield average precinct-level tax incidence per voting-age adult. We detail these procedures in appendix H, in which we argue that our calculations for precinct-level tax incidence are likely to be quite accurate. Of course, an individual household's tax incidence could be higher or lower than the precinct average. But this is not a problem, given that we also model average support for the carbon tax at the precinct level. Intuitively, support for the carbon tax will be lower among households with above-average incidence, and higher among households with below-average incidence. These above- and below-average households will tend to cancel, such that average support will be driven by average incidence within the precinct. Thus, we need only concern ourselves with precinct-level prediction error.

Second, we estimate a logistic model to explain voting on the carbon tax as a function of energy tax incidence, controlling for ideology and census covariates. We show in section H.1 of the appendix how to derive this estimating equation from a standard random-utility discrete-choice model of individual voting behavior. Our precinct-level model takes the form:

$$y_i = \alpha \cdot \text{Total Tax}_i + \beta' \text{Ideology}_i + \gamma' \text{Demographics}_i + \epsilon_i, \quad (1)$$

where $y_i = \ln(s_i/(1 - s_i))$ is mean net utility from the referendum passing and s_i is the share voting "yes" in precinct i , Total Tax is mean energy tax incidence in the precinct

32. For gasoline consumption, we use NHTS data from Washington state. For home energy emissions, we use RECS data from Washington's census division and climate zones.

with corresponding coefficient α , Ideology is a vector of precinct-level ideology variables (Republican vote share plus vote shares on the 12 other ballot measures) with corresponding coefficient vector β , Demographics is a vector of all remaining precinct-level control variables (census demographics and county dummies) with corresponding coefficient vector γ , and ϵ is a precinct-level error. We exclude as controls the census variables we use to calculate tax incidence, including number of vehicles and number of rooms, since tax incidence is a linear function of these variables. Thus, we assume that these variables only correlate with voting via their relationship to tax incidence; that is, tax incidence is a sufficient statistic for these variables. Our remaining, fine-grained controls for age, income, industry, and other factors implicitly capture WTP for other policy impacts, including lower sales taxes, EITC rebates, and shifts in production caused by higher energy costs. We estimate this model in a pooled OLS regression in which we stack both years to yield a single coefficient α on tax incidence but allow coefficients to differ by year for all other variables (including the intercept and county dummies). Our identification assumption is the following: tax incidence is uncorrelated with the error term in our model (i.e., no correlated omitted variables). In these regressions, we weight precinct-years by the total number of votes cast for or against a carbon tax. We interpret the coefficient α as the marginal utility of income, which allows us to rescale and express the dependent variable as WTP in dollars (see sec. H.1 of the appendix).

Third and finally, we use the coefficient estimates from our logistic regression to back out mean WTP for the policy in each precinct. Given our model, total WTP in precinct i is given by $y_i/\hat{\alpha}$, that is, mean net utility from the referendum passing scaled by the estimated marginal utility of income (see the top level of fig. 6). This calculation hinges on a causal interpretation for the estimated coefficient on tax incidence ($\hat{\alpha}$), and immediately yields a decomposition of total WTP into energy tax incidence and other policy impacts (see the middle level of fig. 6). To further decompose WTP for these other policy impacts into components (see the bottom level of fig. 6), we need a stronger assumption: no omitted variables correlated with tax incidence, the ideology variables taken as a whole ($\hat{\beta}'\text{Ideology}_i$), or the demographic variables taken as a whole ($\hat{\gamma}'\text{Other}_i$). Under this stronger assumption, our scaled logistic regression decomposes WTP into energy tax incidence (Total Tax_i) and three additional components: ideology ($\hat{\beta}'\text{Ideology}_i/\hat{\alpha}$), demographics ($\hat{\gamma}'\text{Demographics}_i/\hat{\alpha}$), and a residual ($\hat{\epsilon}/\hat{\alpha}$). To facilitate interpretation, we recenter the ideological variables to equal zero at their sample mean vote shares.³³ Thus, we implicitly measure each additional component of WTP relative to a precinct with average ideology. Below, we report sample means and standard deviations for total

33. That is, we use recentered variables $\tilde{v}_i = v_i - \bar{v}$, where $\bar{v}$ is the sample mean of ideology variable v . Our ideology variables include the Republican vote share in the 2016 presidential election, plus the vote shares on the 12 other ballot measures from 2016 and 2018.

WTP and its various components across precincts (weighted by total votes cast for or against a carbon tax) alongside our coefficient estimates.

4.2. Regression Results and Implied WTP Estimates

Table 5 presents the results from this analysis. We begin by testing the restriction that the coefficient on tax incidence is identical across energy uses and years. Note that Total Tax is the sum of two variables: tax incidence due to gasoline consumption (Vehicle Tax) and tax incidence due to home energy consumption (Room Tax). Column 1 tests that the coefficients on these variables are identical by including Room Tax in addition to Total Tax. The coefficient on Room Tax—equal to the difference in coefficients on Vehicle Tax and Room Tax were we to use these variables instead—is statistically insignificant, implying that we cannot reject identical coefficients on the two underlying sources of tax incidence. Meanwhile, column 2 includes tax incidence interacted with a 2018 dummy (Total Tax 2018), in addition to the variable pooling tax incidence for both years (Total Tax). Again, the coefficient on Total Tax 2018 is statistically insignificant, implying that we cannot reject identical coefficients for the two years. Thus, the remaining columns all impose a single coefficient on tax incidence that is identical across years. Columns 3–4 present results from a single OLS regression. We report the results in two columns (repeating the coefficient estimate of -0.98) so that we can show estimated WTP separately for 2016 and 2018 at the bottom of the table. The coefficient of -0.98 on Total Tax is negative and statistically significant, suggesting that voters derive negative utility from higher taxes, all else equal. Note that we report coefficients on tax incidence measured in \$1,000s (to minimize leading zeros). Thus, the estimated marginal utility of dollars is $\hat{\alpha} = 0.00098$. The coefficient on Total Tax is similar to the coefficient in column 1, which is consistent with most variation in Total Tax coming from Vehicle Tax (see sec. H.2.3 of the appendix). Columns 5–6 likewise present results from a single regression, but estimated via two-stage least squares (2SLS), while columns 7–8 present results from a single regression estimated via OLS with precinct-level fixed effects. We return to these regressions in section 4.3.

We now turn to our estimates of precinct-level WTP reported at the bottom of table 5, focusing on the OLS results in columns 3–4. We report both mean values and standard deviations (below, in brackets). We first focus on mean values. Total WTP is negative \$409 per voter in 2016 and negative \$293 per voter in 2018. Recall that total WTP in a precinct is given by the log-odds ratio scaled by the marginal utility of dollars ($y_i/\hat{\alpha}$) and note that this value equals zero for a precinct whose voters are indifferent on average such that exactly 50% of them would vote yes (since $y_i = \ln(s_i/(1 - s_i)) = 0$ when $s_i = 0.5$). Thus, this value ($y_i/\hat{\alpha}$) intuitively tells us how far in dollars the precinct is from being indifferent to the passage of the referendum. When voters are more sensitive to tax incidence—that is, when α is bigger in magnitude—voters need less monetary compensation to become indifferent, and

Table 5. Logistic Model of Carbon Tax Vote Shares and Implied Willingness-to-Pay Estimates

	OLS (1)	OLS (2)	OLS-16 (3)	OLS-18 (4)	2SLS-16 (5)	2SLS-18 (6)	FE-16 (7)	FE-18 (9)
Total Tax	-1.10* (.50)	-1.06** (.36)	-.98* (.41)	-.98* (.41)	-1.59** (.56)	-1.59** (.56)	-1.49* (.64)	-1.49* (.64)
Room Tax	.97 (.97)							
Total Tax 2018		.31 (.35)						
Total WTP			-409 [554]	-293 [797]	-252 [342]	-181 [492]	-269 [365]	-193 [525]
Energy tax incidence			-135 [22]	-81 [13]	-135 [22]	-81 [13]	-135 [22]	-81 [13]
Other policy impacts			-274 [540]	-212 [789]	-117 [328]	-100 [484]	-134 [351]	-112 [517]
Ideology			-1 [492]	4 [768]	-1 [301]	2 [473]		

Demographics	-273 [70]	-216 [76]	-117 [43]	-102 [47]	...	...
Residual	0 [126]	0 [127]	0 [78]	0 [79]	...	...

Source. Data from Washington Secretary of State, US Census, 2017 National Household Transportation Survey, and 2015 Residential Energy Consumption Survey.

Note. The top part of this table presents coefficient estimates from logistic regressions modeling the share voting “yes” for the carbon tax in 2016 and 2018 as a function of energy tax incidence and various controls. The dependent variable is the log-odds ratio: $y_{it} = \ln(s_{it}/(1 - s_{it}))$, where s_{it} is the share voting “yes” in precinct i in year t . Total Tax measures total tax incidence (in \$1,000s) and is the sum of two variables: Vehicle Tax (gasoline tax incidence as predicted by no. of vehicles, commute time, household size, and urban status) and Room Tax (home energy tax incidence as predicted by no. of rooms, household size, and owner-occupied status). Total Tax 2018 in col. 2 interacts Total Tax with a 2018 dummy. Calculations of tax incidence assume a carbon tax of \$25/tCO₂ in 2016 and \$15/tCO₂ in 2018. All regressions pool data for 2016 and 2018 and include a full set of controls. Controls for ideology are the share voting Republican (vs. Democrat) in the 2016 presidential election, plus the full set of “yes” vote shares on the 12 other ballot measures from 2016 and 2018. Controls for other nontax variables are detailed census variables (industry, home value, income, gender, age, race, and education) and county fixed effects. These pooled regressions also include a dummy variable for the 2018 vote, as well as its interaction with all control variables. For all models, we weight precinct-years by the total votes cast for or against the carbon tax. The regressions in cols. 1, 2, and 3–4 are estimated via OLS. The regression in cols. 5–6 is estimated via 2SLS. Instruments include the following variables, plus their interactions with a 2018 dummy, all divided by the no. of voting-age adults per household: share of households with 1+ vehicles, commute time, household size, share urban, no. of rooms, share owner occupied, and a constant term. The regression in cols. 7–8 is estimated via OLS and includes precinct-level fixed effects, such that coefficients are identified only for variables that differ over time within precinct: Total Tax, the 2018 dummy, and the 2018 dummy interacted with control variables. Standard errors are clustered by county (39 clusters). The bottom part of the table reports precinct-level means for estimated WTP (in \$), with the corresponding standard deviations appearing immediately below (in brackets), after weighting by total votes cast for or against the carbon tax. The first row reports total WTP. The next two rows report the decomposition of total WTP into energy tax incidence and other policy impacts. The last three rows report the decomposition of other policy impacts into components explained by ideology, demographics, and a residual. Columns 3–4 are based on a single pooled regression and likewise for cols. 5–6 and 7–8; we use two columns to report WTP estimates separately by year.

- * Significant at the 5% level.
- ** Significant at the 1% level.
- *** Significant at the .1% level.

therefore the mean WTP for the policy is smaller in magnitude. Of course, these results hinge on a causal interpretation for the coefficient on Total Tax.

How much of this negative WTP is due to energy tax incidence? Energy tax incidence accounts for negative \$135 in 2016 (col. 3) and \$81 in 2018 (col. 4). By construction, these values equal the average energy tax incidence across precincts in these years. This leaves other policy impacts (such as sales tax rebates, EITC, green projects, and lower emissions) accounting for negative \$274 in 2016 (col. 3) and negative \$212 (col. 4). That is, even if the energy tax incidence were zero, voters would have rejected these referendums. Apparently, they do not like these other policy impacts, taken as a whole. Alternatively, they do like these other policy impacts but overestimate the costs of higher energy taxes (Douenne and Fabre 2022).

If we further assume a causal interpretation for our two sets of control variables, then we can decompose WTP for these other policy impacts more fully into components attributable to ideology, demographics, and a residual. In this case, mean WTP attributable to ideology is approximately zero—a direct consequence of our choice to recenter ideology to equal zero on average. Finally, mean WTP attributable to demographics is negative \$273 in 2016 (col. 3) and negative \$216 in 2018 (col. 4). Note that our decomposition implicitly measures mean WTP attributable to demographics relative to a precinct with average ideology. Of course, the residual component is zero on average—a numerical property of OLS regression.

Overall, we find a negative WTP for the carbon tax policy between \$293 and \$409. Looking at the determinants of WTP, we find that the average voter (precinct, strictly speaking) dislikes the personal cost of higher energy taxes, as expected—but seems barely to credit the potential benefits of a carbon tax. These results contrast with the stated-preference estimates of Kotchen et al. (2017), who find a positive WTP of \$177 on average for the nontax benefits of a carbon tax.³⁴ Proponents of the 2016 initiative hoped that returning carbon tax revenue to residents through a lower state sales tax and expanded EITC would directly counteract voters' dislike for the carbon tax, while proponents of the 2018 initiative hoped to boost support through spending on green projects. However, we estimate a WTP of negative \$212 to \$274 for these other policy impacts, suggesting that both strategies failed. The results are consistent with voters thinking that the benefits would not reach them. Such lack of trust in public authorities is one of the issues believed to hinder the adoption of carbon taxes (Carattini et al. 2018).

To understand what drives variation in WTP across precincts, we now turn to the standard deviations (in brackets). Note that the standard deviations due to energy tax incidence (\$13–\$22), demographic variables (\$70–\$76), and residual (\$126–\$127)

34. Note that Kotchen et al. (2017) estimate average preferences nationwide for a national-level carbon tax, whereas we estimate average preferences in Washington for a state-level carbon tax.

are all swamped by the standard deviation in WTP due to ideology (\$492–\$768), which is consistent with our results above that ideology explains most of the variation in voting across precincts. These results underscore the importance of factors other than self-interest in shaping policy preferences (Luttmer 2001).

Figure 7 illustrates these points graphically with estimated marginal WTP curves for the two carbon tax policies, that is, mean WTP as we move through the distribution of precincts sorted from high to low support. These figures are based on the regression results reported in columns 3 and 4 of table 5 and group precincts by percentile of WTP. We show both total WTP for the marginal precincts assuming a causal interpretation for the coefficient on tax incidence as well as a breakdown of this WTP into individual components (energy tax incidence, ideology, demographics, and residual) assuming causal interpretations for ideology and other control variables. The WTP attributable to taxes, which averages –\$135 in 2016 and –\$81 in 2018, is nearly constant throughout the distribution in both years. This finding implies that the characteristics that drive tax incidence (mainly vehicle ownership and home size) do not differ substantially across precincts with high versus low support for the carbon tax. The WTP attributable to demographics is also fairly constant across precincts. The residual component is \$0 on average but turns slightly positive in the left tail and slightly negative in the right: unobservables tend to be relatively more important in explaining extreme voting outcomes than average outcomes. Ideology is clearly the main driver of variation in WTP across precincts, ranging from about –\$1,100 to \$1,200 in 2016 and about –\$1,400 to \$1,900 in 2018. The substantially wider range in 2018 is consistent with our findings above that the 2018 policy was more ideologically slanted, gaining support in liberal precincts and losing support in conservative ones. Finally, note that if we focus on the median precinct (0.5 on the horizontal axis) and look closely, we see that WTP due to the tax accounts for about one-quarter to one-third of overall negative WTP, which is consistent with the summary statistics at the bottom of table 5.

4.3. Robustness of Our Willingness-to-Pay Estimates

In this section, we explore the robustness of our WTP estimates. We first consider and rule out the possibility that individual-level variation in energy tax incidence is large enough to overturn our conclusion that ideology trumps tax incidence. Then, we explore the sensitivity of our results to changes in the coefficient on Total Tax, differentiating by the cause of this change: behavioral issues, omitted variables bias, and measurement of energy tax incidence.

4.3.1. Individual-Level Variation in Tax Incidence

We estimate the mean and standard deviation in WTP across precincts and report the results in table 5. Based on this analysis, we conclude that variation in precinct-level WTP is mainly driven by ideology, rather than by tax incidence. Of course, this analysis ignores within-precinct variation in WTP. In principle, if within-precinct variation




Figure 7. Marginal willingness to pay for a carbon tax by precinct. This figure plots marginal WTP curves for the 2016 (*top panel*) and 2018 (*bottom panel*) carbon tax policies. The horizontal axis ranks precincts from highest to lowest total WTP for a carbon tax in each year. The vertical axis measures, for the marginal precinct, total WTP along with a breakdown of its individual components (tax incidence, ideology, other, and residual). To estimate marginal WTP, we sort precincts from lowest to highest WTP, divide them into percentiles, and calculate mean WTP within each percentile. In constructing percentiles, we weight by total carbon tax votes cast; thus, each percentile represents the same number of voters. Likewise, when calculating mean WTP within each percentile, we weight by total votes cast. Data source: Washington Secretary of State, US Census, 2017 National Household Transportation Survey, and 2015 Residential Energy Consumption Survey.

in tax incidence were large relative to within-precinct variation in ideology, then this could overturn our conclusion that ideology trumps incidence (we thank an anonymous referee for clarifying this point). To address this concern, we present a bounding exercise in section H.3 of the appendix, which shows that the standard deviation of tax incidence across all individuals ranges \$77–\$100 in 2016 and \$46–\$60 in 2018. These values exceed the estimated standard deviation in tax incidence across precincts: \$22 in 2016 and \$13 in 2018 (see table 5). But they are still much lower than the standard deviation in WTP across precincts that can be attributed to ideology: \$492 in 2016 and \$768 in 2018 (see cols. 3–4 in table 5). Note that these latter values are themselves lower bounds on the standard deviation in WTP across individuals that can be attributed to ideology, since they ignore within-precinct variation. Thus, we still conclude that ideology trumps tax incidence, both for precincts on average and for individuals.

4.3.2. Behavioral Factors

Figure 5 suggests that votes depend on a binary “car versus no car” distinction, rather than the continuous number of vehicles. If voters use a binary heuristic, then our measure of tax incidence will diverge from their beliefs, leading to bias. To account for this divergence, we purge any variation in Total Tax coming from the continuous number of vehicles. Mechanically, we do this by estimating our model via 2SLS, instrumenting for Total Tax with percentage of households with 1+ vehicles (the binary measure), plus the other continuous variables used above to estimate tax incidence.³⁵ In the first stage, we regress Total Tax on the instruments and controls. In the second stage, we regress support for the carbon tax on predicted tax incidence and controls. Thus, our assumption is essentially the same as above: the excluded instruments used to predict tax incidence only correlate with voting via their relationship to tax incidence. Importantly, we are not using 2SLS to overcome any fundamental endogeneity concern present in our OLS approach, such as omitted variables. Rather, we are using 2SLS purely for numerical reasons: to purge variation in tax incidence coming from the continuous number of vehicles (which fig. 5 indicates is not salient to voters) while preserving all other sources of variation used to construct this measure.

Columns 5–6 of table 5 present the results from this 2SLS regression. As expected, when we focus on variation in Total Tax driven by the binary measure of vehicle

35. Our full set of instruments includes percentage of households with 1+ vehicles, commute time, household size, percentage urban, number of rooms, percentage owner-occupied, and a constant term, plus interactions of these variables with a 2018 dummy (to capture the lower carbon price in that year), all divided by the number of voting-age adults per household (given our rescaling of household-level tax incidence per voting-age adult). Note that these instruments would perfectly predict Total Tax were we to use the average number of vehicles per household instead of the percentage of households with 1+ vehicles.

ownership, the coefficient on Total Tax increases in magnitude to -1.59 . Thus, mean total WTP at the bottom of the table shrinks proportionately to $-\$252$ in 2016 and $-\$181$ in 2018 (remember that WTP is $y_i/\hat{\alpha}$, where y_i is the transformed vote share, and to divide the table coefficients by 1,000 to get $\hat{\alpha}$), while mean tax incidence remains unchanged in both 2016 and 2018 (by construction). Thus, after accounting for behavioral factors, tax incidence is relatively more important in explaining overall negative WTP for the policies. Meanwhile, mean WTP attributable to ideology remains near zero (by construction, due to our normalization). Thus, given the overall increase in mean WTP (becoming less negative), the WTP attributable to the “other” category must increase in tandem (becoming less negative in cols. 5–6). These results show that our estimate of overall WTP is somewhat sensitive to specification, as is the amount that we attribute to nontax factors. However, qualitatively, the variation in WTP across precincts (standard deviation of $\$342$ – $\$492$) continues to be dominated by variation that we attribute to ideology (standard deviation $\$301$ – $\$473$).

4.3.3. Omitted Variables Bias

Omitted variables bias in the coefficient on Total Tax has a qualitatively similar effect as that described above. If the true coefficient increases in magnitude, mean WTP shrinks and we attribute a larger share to tax incidence. Meanwhile, if the coefficient decreases in magnitude, these effects are reversed. How sensitive are our qualitative results to the coefficient on Total Tax? Note that total WTP is three times larger than mean tax incidence in column 3, while the ratio is somewhat larger in column 4. Thus, the true coefficient on Total Tax would need to be at least 1.5 times larger for mean tax incidence to account for a majority of the negative WTP overall and three times larger for mean tax incidence to exactly equal mean WTP overall. Any larger increase would imply that mean WTP attributable to nontax factors is strictly positive. Meanwhile, note that the standard deviation of total WTP is 25 times larger than the standard deviation of tax incidence in column 3, and the ratio is even larger in column 4. Thus, the true coefficient on Total Tax would need to be at least 25 times larger to overturn our conclusion that variance in total WTP across precincts is dominated by ideology and other nontax factors.

We partially address the omitted variables concern by adding precinct-level fixed effects to our pooled OLS regression. Fixed effects control for time-constant unobservables that might be correlated with carbon footprint across precincts. Variation in tax incidence comes from the decline in tax rate from $\$25/\text{tCO}_2$ in 2016 to $\$15/\text{tCO}_2$ in 2018, interacting with variation in carbon footprint across precincts. This is essentially a difference-in-differences estimator comparing precincts with high CO_2 emissions to precincts with low CO_2 emissions before versus after the change in tax rate. Recall that the 2016 and 2018 policies also differed in terms of revenue recycling. Thus, we continue to control for ideology, census demographics, and county fixed effects, all interacted with the 2018 dummy (the controls themselves do not vary over

time and therefore drop out of this estimation). Thus, our identification assumption is that changes in tax incidence are uncorrelated with unobserved factors that determine changes in support, conditional on our myriad control variables.

Columns 7–8 of table 5 report the results of this fixed-effects (FE) regression. The coefficient on total tax is 1.5 times larger than in the pooled OLS regression in columns 3–4. Thus, energy tax incidence now accounts for a larger share of the overall negative WTP (about half). But the variation in total WTP across precincts (standard deviation of \$365–\$525) is still dominated by other policy impacts (standard deviation of \$351–\$517). Our inclusion of precinct-level fixed effects prevents us from decomposing WTP further in these regressions since we can only identify how ideology and demographics relate to changes in WTP. But consistent with our earlier findings, auxiliary regressions show that ideology explains 92% of the variation in WTP not explained by taxes in 2016 and 96% of this variation in 2018.

4.3.4. *Measurement of Energy Tax Incidence*

We assume that voters perceive carbon prices of \$25/tCO₂ in 2016 and \$15/tCO₂ in 2018, since these are the values used by tax calculators available to Washington voters at the time. However, voters may have perceived higher overall prices based on the phased-in tax increases under both policies. See section H.2.3 of the appendix for an extended discussion. Suppose we doubled the taxes perceived by voters in both years (not unreasonable, based on our reading of the legislation and contemporaneous reporting). This rescaling would increase precinct-level tax incidence by a factor of two. Meanwhile, the coefficient on Total Tax in table 5 would mechanically shrink by a factor of two, which would inflate all of our other WTP estimates by the same factor. Thus, everything WTP-related in our reported results would get multiplied by two, including the means and standard deviations reported at the bottom of table 5, along with the number labels on the vertical axis in figure 7. Thus, voters perceiving a proportionally higher tax incidence would merely multiply all WTP components by the same factor and thus would not change our conclusions about the relative role of the various components of WTP.

5. COULD A CARBON TAX PASS IN OTHER STATES?

How would policies like I-732 and I-1631 fare in other states? This is a key question for carbon tax policy entrepreneurs hoping to apply lessons from the Washington experience to other states. To address it, we construct a set of 50 out-of-sample forecasts by applying coefficients from a regression of the precinct-level vote in Washington State to the observed state-level demographics and presidential vote shares in all 50 states.³⁶ We

36. Our state-level demographics come from the US Census and measure the same variables in the same years as our precinct-level data from Washington but at the state level (e.g., share of

base these forecasts on the same specification as model 3 in table 3 (Republican share plus census demographics) but estimated separately for 2016 and 2018. See appendix F for these results. In principle, we could improve these forecasts by controlling for ideology, as in model 4, but we do not observe this measure for other states (since our measure relies on Washington-specific ballot initiatives). In practice, there would be little gain in precision: simple regressions of the predicted values from model 4 on the predicted values from model 3 yield coefficients of 1.000 and R -squared values of 0.97–0.98.

Technically, these 50 out-of-sample forecasts tell us how Washington would have voted differently on I-732 and I-1631 if we reweighted its precincts to have the same demographics and ideology on average as other states. To go further and interpret these forecasts as predictions for how other states would have voted, we must assume that the coefficients that relate voting on the carbon tax to ideology and demographics in Washington would transfer to other states. Of course, we cannot take this latter assumption too literally. First, and most obviously, many states do not even feature a popular initiative mechanism. Second, measures identical to I-732 and I-1631 would be impossible in many states because they lack a state sales tax or have restrictions on how various sources of tax revenue may be used. Third, the incidence of a carbon tax would certainly differ across states due to differences in demand for home heating and cooling, energy sources used for electricity generation, and preexisting regulations such as California's cap-and-trade program. Finally, the information environment that exists in Washington generally given its 100% mail-in ballots and universal distribution of voter guides, or that existed in 2016 and 2018 given the events described above, would not transfer exactly to other times and locations. Thus, our forecasts are merely illustrative: they demonstrate the relative challenges facing carbon tax advocates in various states, but we cannot take the specific quantitative results literally.

Figure 8 plots the resulting state-level forecasts for “yes” on I-732 (hollow dots) and I-1631 (full dots) versus the ranked Republican vote share in the 2016 presidential election. States that feature the popular initiative are shown in black, while states that lack a popular initiative are shown in gray. Overall, states with higher Republican vote shares would be less likely to pass these measures, which is consistent with ideology driving the vote for a carbon tax. The individual state forecasts do not decrease

people in California aged 40–44). Our presidential voting data come from the US Federal Elections Commission and record statewide vote shares for various parties in the 2016 US presidential election. See <https://transition.fec.gov/pubrec/fe2016/federalelections2016.pdf>. We use these data to calculate the share voting Republican (vs. Democrat) for each state. Recall that we define our variables based on shares of voters, people, households, and so forth living in different Washington precincts. Thus, we can apply our regression coefficients directly to the corresponding state-level aggregate variables from other states and obtain the same forecast were we to instead apply our coefficients to the corresponding precinct-level variables from other states (which we do not observe) and then aggregate to the state level.




Figure 8. Forecast carbon tax vote share by state, for the 2016 and 2018 carbon tax versions. This figure plots out-of-sample forecasts by state for the share voting “yes” on I-732 in 2016 (hollow dots) and I-1631 in 2018 (full dots) versus the ranked Republican vote share in the 2016 presidential election (with connecting lines to facilitate comparison). Forecasts are based on estimating model 3 in table 3 but with separate regressions for 2016 and 2018, using precinct-level data from Washington State. See appendix F. State forecasts are generated by applying coefficients from these precinct-level regressions to the same variables measured at the state level in all 50 states. States that feature a popular initiative mechanism are shown in black, while states that lack an initiative process are shown in gray. Data source: Washington Secretary of State, US Census, and US Federal Election Commission.

monotonically with Republican vote share due to variation in the composition of the electorate for variables other than the Republican vote share.

Which states would be most likely to pass a carbon tax? Figure 9 zooms in on the 10 states with the lowest Republican share and provides confidence intervals. Only three of these states—Washington, Massachusetts, and California—feature a popular initiative mechanism.³⁷ Even ignoring that fact, we forecast that no other state would pass I-732, while just one state—Vermont, which lacks a popular initiative—would pass the more progressive I-1631. Of the initiative states, we forecast that Massachusetts comes closest to passing I-1631 with 49.1% of votes in favor and a 95% confidence interval that reaches to 49.7%.

37. Moreover, the difficulty in getting on the ballot is highly variable across states. Massachusetts is notoriously difficult to navigate, while Washington and California are easier.




Figure 9. Top 10 most Democratic states: forecast carbon tax vote share by state for the 2016 and 2018 carbon taxes. This figure plots out-of-sample forecasts by state for the share voting “yes” on I-732 in 2016 (hollow dots) and I-1631 in 2018 (full dots) versus the ranked Republican vote share in the 2016 presidential election. Forecasts are based on estimating model 3 in table 3 but with separate regressions for 2016 and 2018, as reported in appendix F, using precinct-level data from Washington State. State forecasts are generated by applying coefficients from these precinct-level regressions to the same variables measured at the state level in all 50 states. States that feature a popular initiative mechanism are shown in black, while states that lack an initiative process are shown in gray. Point estimates are represented by circles, while 95% confidence intervals are represented by vertical lines. Data source: Washington Secretary of State, US Census, and US Federal Election Commission.

So then: How liberal would a state need to be to pass a carbon tax? Take Washington. Forecast support there is 41.7% in 2016 (8.3% below passing), while the coefficient on Republican share for that year is -0.523 (see app. F). Thus, Republican share would need to be $8.3\% / 0.523 \approx 16\%$ lower—holding demographics fixed—to pass I-732. Likewise, Washington’s forecast support is 43.5% in 2018 (6.5% below passing), while the coefficient on Republican share is -0.747 (see app. F). Thus, Republican share would need to be $6.5\% / 0.747 \approx 8.7\%$ lower to pass I-1631. Similar calculations can be performed for other states: just multiply the passage gap (50% vote share minus our forecast) by $1/0.523 \approx 1.9$ in 2016 and $1/0.747 \approx 1.3$ in 2018.

Since I-1631 is more progressive than I-732, we should expect it to perform relatively better in liberal states and relatively worse in conservative states. This is illustrated in figure 8. On the left side of the figure (more liberal), the I-1631 forecast (full

dots) tends to exceed the I-732 forecast (hollow dots). Meanwhile, on the right side of the figure (more conservative), the pattern is reversed.³⁸ However, there is still important variation across states explained by other voter characteristics. For example, among the 10 most liberal states in figure 9, we forecast that I-1631 would do better than I-732 in Massachusetts and Vermont but marginally worse in California and Hawaii.³⁹

Generally, states for which we forecast a higher vote share in figure 8 already regulate carbon, though not through a carbon tax. Many of these states have a cap-and-trade mechanism, suggesting that cap and trade might be easier to pass politically than a carbon tax. In particular, California has an economy-wide cap-and-trade program, while many northeastern states participate in the Regional Greenhouse Gas Initiative (RGGI) and its cap-and-trade program for electricity-sector emissions. These states include Vermont (highest forecasted vote share for I-1631), Massachusetts, New York, Maryland, Rhode Island, Connecticut, Delaware, Maine, and New Hampshire. Meanwhile, Washington itself passed cap and trade in April 2021—not via referendum but rather via the legislature (SB 5126). Governor Jay Inslee signed the bill in May 2021. As another indicator, many states with high forecast vote shares are members of the US Climate Alliance—a coalition of states that has committed to meeting the Paris Climate Accord’s abatement goals. This alliance includes California, six of nine RGGI states, and Hawaii, Washington, Oregon, Colorado, Virginia, Minnesota, and North Carolina.⁴⁰

6. CONCLUSION

Climate policy is one of the most politically polarized topics in the United States, making the adoption of federal legislation to limit greenhouse emissions especially difficult. States represent a potentially valuable laboratory for learning about the politics of climate policy—including the politics of Pigouvian taxes on carbon emissions, which are the subject of substantial research by economists.

We analyze two failed carbon tax initiatives in Washington State using precinct-level aggregate voting data from Washington. We estimate that the direct economic

38. We confirm this pattern in app. I, which directly forecasts the difference in state vote shares based on model 3 in table 4.

39. Our forecast that I-1631 performs marginally worse in California and Hawaii is due, in part, to the high shares of Asians and Pacific Islanders in these states—and the fact that, within Washington, I-1631 tends to perform worse than I-732 in precincts with higher shares for these demographic groups. Indeed, when we omit our controls for race and ethnicity altogether, we forecast that I-1631 performs better than I-732 in California and Hawaii.

40. For a list of members, a statement of principles, and further background, see <https://www.usclimatealliance.org/>. Maryland, Maine, and New Hampshire are the only RGGI states that are not alliance members—and note that Maine and New Hampshire have the lowest forecast vote shares among RGGI states.

incidence of the carbon tax is similar across Washington's precincts on average. Thus, for the median precinct, resistance to higher energy prices is an important factor to explain why voters rejected the two carbon taxes. However, ideology can easily overpower pocketbook concerns. Indeed, we show that ideology is by far the best and most important predictor of variation in support for the carbon tax across precincts, explaining more than 90% of the variation in vote shares.

What do our results imply about the prospects for carbon taxes at the state level? After all, nationwide surveys show strong majorities in favor and support that extends even to the most conservative states in the country. However, we present evidence based on actual voting showing that this support is rather shallow and easily overturned with a vigorous (and often negative) political campaign. Overall, our results suggest that a carbon tax modeled on recent policies from Washington State would—without deeper efforts to inform and persuade voters—be unlikely to pass in the near term. The best prospects for a carbon tax continue to be in liberal states, especially in Massachusetts. Furthermore, legislation and unilateral executive action by governors (like Inslee's line-item veto of hold-ups to cap and trade) would appear to be more promising avenues to experiment with state carbon taxes.

What political levers might carbon tax advocates pull to boost support and deflect opposition? Washington's experience shows that a revenue-neutral policy can appeal to moderates and conservatives—or at least limit their opposition—while a tax-and-spend policy can appeal to liberals. But careful policy design does not seem to be enough, for policies aimed at both ends of the political spectrum were tried in Washington State, and both failed. For a tax-and-spend policy like I-1631 to succeed at the polls, voters must be convinced that the new spending has real value. For a revenue-neutral policy like I-732 to succeed, voters must be informed and convinced that—even though their energy bills will go up—they will benefit from lower taxes elsewhere. A carbon tax-and-dividend policy, whose cash benefits may be more transparent to voters, is a promising avenue for future research.

REFERENCES

- Amdur, David, Barry G. Rabe, and Christopher Borick. 2014. Public views on a carbon tax depend on the proposed use of revenue. *Issues in Energy and Environmental Policy*, no. 13.
- Anderson, Soren T., Ryan Kellogg, and James M. Sallee. 2013. What do consumers believe about future gasoline prices? *Journal of Environmental Economics and Management* 66 (3): 383–403.
- Anderson, Soren T., and James M. Sallee. 2016. Designing policies to make cars greener. *Annual Review of Resource Economics* 8 (1): 157–80.
- Bornstein, Nicholas, and Bruno Lanz. 2008. Voting on the environment: Price or ideology? Evidence from Swiss referendums. *Ecological Economics* 67 (3): 430–40.
- Burkhardt, Jesse, and Nathan Chan. 2017. The dollars and sense of ballot propositions: Estimating willingness to pay for public goods using aggregate voting data. *Journal of the Association of Environmental and Resource Economists* 4 (2): 479–503.

- Carattini, Stefano, Andrea Baranzini, Philippe Thalmann, Frédéric Varone, and Frank Vöhringer. 2017. Green taxes in a post-Paris world: Are millions of nays inevitable? *Environmental and Resource Economics* 68 (1): 97–128.
- Carattini, Stefano, Maria Carvalho, and Sam Fankhauser. 2018. Overcoming public resistance to carbon taxes. *Wiley Interdisciplinary Reviews: Climate Change* 9 (5): e531.
- Converse, Phillip E. 1964. The nature of belief systems in mass publics. In *Ideology and discontent*, ed. David E. Apter. New York: Free Press.
- Douenne, Thomas, and Adrien Fabre. 2022. Yellow vests, pessimistic beliefs, and carbon tax aversion. *American Economic Journal: Economic Policy* 14 (1): 81–110.
- Egan, Patrick J., and Megan Mullin. 2017. Climate change: U.S. public opinion. *Annual Review of Political Science* 20 (1): 209–27.
- Fujita, Masahisa. 1989. *Urban economic theory: Land use and city size*. Cambridge: Cambridge University Press.
- Gerber, Elisabeth R., and Jeffrey B. Lewis. 2004. Beyond the median: Voter preferences, district heterogeneity, and political representation. *Journal of Political Economy* 112 (6): 1364–83.
- Holian, Matthew J., and Matthew E. Kahn. 2015. Household demand for low carbon policies: Evidence from California. *Journal of the Association of Environmental and Resource Economists* 2 (2): 205–34.
- Kahn, Matthew E., and John G. Matsusaka. 1997. Demand for environmental goods: Evidence from voting patterns on California initiatives. *Journal of Law and Economics* 40 (1): 137–74.
- Knittel, Christopher R. 2014. The political economy of gasoline taxes: Lessons from the oil embargo. *Tax Policy and the Economy* 28 (1): 97–131.
- Kotchen, Matthew J., Zachary M. Turk, and Anthony A. Leiserowitz. 2017. Public willingness to pay for a U.S. carbon tax and preferences for spending the revenue. *Environmental Research Letters* 12 (9): 094012.
- Lang, Corey, and Shanna Pearson-Merkowitz. 2021. Aggregate data yields biased estimates of voter preferences. Working paper (March).
- Luttmer, Erzo F. P. 2001. Group loyalty and the taste for redistribution. *Journal of Political Economy* 109 (3): 500–528.
- Masket, Seth E, and Hans Noel. 2012. Serving two masters: Using referenda to assess partisan versus dyadic legislative representation. *Political Research Quarterly* 65 (1): 104–23.
- Mildenberger, Matto, Erick Lachapelle, Kathryn Harrison, and Isabelle Stadelmann-Steffen. 2022. Limited impacts of carbon tax rebate programmes on public support for carbon pricing. *Nature Climate Change* 12:141–47.
- Rabe, Barry G. 2018. *Can we price carbon?* Cambridge, MA: MIT Press.
- Snyder, James M., Jr. 1996. Constituency preferences: California ballot propositions, 1974–90. *Legislative Studies Quarterly* 21 (4): 463–88.
- Thalmann, Philippe. 2004. The public acceptance of green taxes: 2 million voters express their opinion. *Public Choice* 119 (1): 179–217.
- von Thünen, J. H. 1826. *Der isolierte Staat in Beziehung auf Landwirtschaft und Nationalökonomie*. Hamburg: Perthes. English translation by C. M. Wartenberg, *Von Thünen's Isolated State*. Oxford, Pergammon Press, 1966.
- Wooldridge, Jeffrey M. 2020. *Introductory econometrics: A modern approach*. Boston: Cengage.