NBER WORKING PAPER SERIES

WHERE DO MY TAX DOLLARS GO? TAX MORALE EFFECTS OF PERCEIVED GOVERNMENT SPENDING

Matias Giaccobasso
Brad C. Nathan
Ricardo Perez-Truglia
Alejandro Zentner

Working Paper 29789
http://www.nber.org/papers/w29789

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2022, Revised February 2023

We are thankful for excellent comments from several colleagues. This project was reviewed and approved in advance by the Institutional Review Board at The University of Texas at Dallas. The field experiment was pre-registered in the AEA RCT Registry (#0007483). To prevent contamination of the subject pool (e.g., that subjects or the media could read about the hypotheses), we posted the RCT pre-registration immediately after the deadline to file a protest had passed, but before conducting any analysis of the data. After the study is accepted for publication, we will share all the code and data through a public repository. Xinmei Yang provided superb research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Where Do My Tax Dollars Go? Tax Morale Effects of Perceived Government Spending
Matias Giaccobasso, Brad C. Nathan, Ricardo Perez-Truglia, and Alejandro Zentner
NBER Working Paper No. 29789
February 2022, Revised February 2023
JEL No. C93,H26,I22,Z13

ABSTRACT

Do perceptions about how the government spends tax dollars affect the willingness to pay taxes? We designed a field experiment to test this hypothesis in a natural, high-stakes context and via revealed preferences. We measure perceptions about the share of property tax revenues that fund public schools and the share of property taxes that are redistributed to disadvantaged districts. We find that even though information on where tax dollars go is publicly available and easily accessible, taxpayers still have significant misperceptions. We use an information-provision experiment to induce exogenous shocks to these perceptions. Using administrative data on tax appeals, we measure the causal effect of perceived government spending on the willingness to pay taxes. We find that some perceptions about government spending have a significant effect on the probability of filing a tax appeal and in a manner that is consistent with the classical theory of benefit-based taxation. We discuss implications for researchers and policy makers.

Matias Giaccobasso
Anderson School of Management - UCLA
110 Westwood Plaza, C 3.10
Los Angeles, CA 90077
mgiaccobasso@ad.ucla.edu

Brad C. Nathan
Naveen Jindal School of Management
The University of Texas at Dallas
800 W Campbell Rd.
Richardson, TX 75080
brad.nathan@utdallas.edu

Ricardo Perez-Truglia
Haas School of Business
University of California, Berkeley
545 Student Services Building #1900
Berkeley, CA 94720-1900
and NBER
ricardotruglia@berkeley.edu

Alejandro Zentner
Naveen Jindal School of Management
University of Texas at Dallas
800 W Campbell Rd.
Richardson, TX 75080
azentner@utdallas.edu

A data appendix is available at http://www.nber.org/data-appendix/w29789

1 Introduction

Perceptions about how the government spends tax dollars may affect the willingness to pay taxes. We test this hypothesis using a natural field experiment in a high-stakes context and via revealed-preferences.

Property taxes represent an important source of revenue for governments in the United States and around the world. For instance, U.S. property tax revenues in 2019 were estimated at $577 billion (Tax Policy Center, 2021a).[1] In the United States, virtually all counties rely heavily on property taxes to fund key government services such as schools, parks, and roads. School funding typically makes up the largest component of property taxes.

There are different ways in which taxpayers' perceptions about government spending could affect their willingness to pay taxes. In this paper, we focus on one mechanism that was studied by early tax policy scholars: benefit-based taxation (Seligman, 1908; Musgrave, 1959). The logic of benefit-based taxation is that individuals are willing to pay more taxes when they believe that they benefit from the government services funded via those taxes. Our research design leverages the fact that some households benefit directly from public education because they have children enrolled in local public schools, whereas others do not benefit directly because they have no children enrolled in public schools. For the sake of brevity, we refer to households with children enrolled in public schools as "households *with* children" and those without children enrolled in public schools as "households *without* children."

When studying attitudes towards taxation, economists and other social scientists rely primarily on survey data (Alesina and Giuliano, 2011; Cruces et al., 2013). However, survey data have some well-known limitations, such as social desirability bias. For example, some individuals may *say* that they are willing to pay more in taxes but would *choose* otherwise when facing real stakes. We study the willingness to pay taxes via revealed preferences, using households' decisions to file tax appeals (Nathan et al., 2020). Filing an appeal is a consequential action that households can take to reduce the amount they have to pay in property taxes.

We conducted our field experiment in Dallas County, Texas. We focus on this county primarily because, from a logistical perspective, it is more practical to implement a field experiment in a single location. However, as property taxes and tax appeals work similarly in most of the United States and in other countries around the world (Nathan et al., 2020; Dobay et al., 2019; World Bank, 2019), our experimental design could be easily replicated in other locations. Dallas County is the second-largest county in Texas, with an estimated

[1] For reference, the 2019 federal income tax generated $1.717 trillion in revenue and corporate income tax $230 billion (Tax Policy Center, 2021b).

population of about 2.6 million in 2020 (U.S. Census Bureau, 2021) – indeed, Dallas County has a larger population than 15 of the 50 U.S. states. The county also is diverse along many dimensions such as ethnicity and partisanship.[2] We designed a survey to be completed by homeowners immediately before they face the opportunity to file a tax appeal. The survey elicits whether the household has children enrolled in public schools to identify which subjects benefit directly from public school spending and which do not.

In the first treatment arm, we measure respondents' perceptions about the share of their own property taxes that corresponds to school taxes and thus funds public school spending. For the sake of brevity, in the remainder of the paper we refer to this percentage as the household's "school share." For reference, the school share is around 49.78% for the average household in Dallas County.[3] We can measure the respondents' misperceptions about where their tax dollars go by comparing their guesses about the school share to the true estimates from administrative records. To study the causal effect of beliefs about government spending, the survey embeds an information-provision experiment. After eliciting respondents' prior beliefs, we inform a random half of them about the true value of their respective school shares. By doing so, we can assess how that information affects their posterior beliefs, as measured by our survey, and their decisions to file a tax appeal, as measured by administrative data.

The information-provision experiment creates exogenous variation in respondents' posterior beliefs about the fraction of their property taxes that funds local schools. For example, a subject who perceives her or his school share amount to be 30% may be informed that the actual share is 50%. According to the benefit-based framework and as noted in the randomized control trial (RCT) pre-registration, the expected effects of the information shock depends on whether the household has children enrolled in public schools. Upon learning that the school share is higher than initially thought, households *with* children should become less likely to file a tax appeal, because they learn that they benefit more from government services than they originally believed. Conversely, households *without* children enrolled in public schools should become more likely to file a tax appeal, because they learn that they benefit less from government services than they originally thought.

The logic of benefit-based taxation could have implications for the ability to redistribute taxes. When taxpayers learn that their tax dollars are being spent in communities other than their own, they may be less willing to pay taxes. We explore this additional hypothesis using a second treatment arm and leveraging the significant redistribution of property taxes that

[2] For example, in the 2012 presidential election, Barack Obama received 57% of the votes, whereas Mitt Romney received 42%, and the remaining 1% of votes went to third-party candidates.

[3] This average is calculated over 400,192 properties in Dallas County and excludes properties such as commercial properties and non-owner-occupied residencies; for more details, see Appendix B.2. Unless explicitly stated otherwise, all statistics about Dallas County are based on this sample.

occurs in some states, including Texas. In 1993, Texas enacted legislation to make funding more equitable across school districts by redistributing a significant fraction of property taxes from property-wealthy districts to property-poor districts.[4] The media often refer to this legislation as the "Recapture Plan" or the "Robin Hood Plan." Due to the high stakes involved, the recapture system has been a topic of heated debate among politicians and the general public (Dallas Morning News, 2018).

In the second treatment arm, we measure perceptions about the share of a household's school taxes that are redistributed across school districts. For the sake of brevity, in the remainder of the paper we refer to this as the "recapture share." For example, a recapture share of 50% would imply that half of the school tax revenue will be transferred to disadvantaged school districts. We can measure respondents' misperceptions about the recapture share by comparing their guesses to actual estimates from administrative records. Again, we measure the causal effects of the perceived recapture share using an information-provision experiment. According to the framework of benefit-based taxation, the effects of the recapture share should depend on whether the household has children. For example, households *with* children should be more likely to protest upon learning that their school taxes are being diverted to other districts because these households value and use the government service. In contrast, households *without* children should not be more likely to protest because the diverted funds are being used for a service that do not benefit them directly.

We sent a letter to a sample of households in Dallas County inviting them to participate in an online survey. Our main subject pool comprises 2,110 respondents who completed the survey. We conducted the survey in April and May of 2021, a period during which subjects could file a protest of their property taxes with the county.[5] We then matched survey responses to administrative records from the county assessor's office. The rich administrative data allowed us to determine, among many other things, if the survey respondent subsequently filed a tax appeal.

The average subject in our sample owns a home worth $349,988 and pays $7,738, on average, in annual property taxes.These estimated taxes are prior to the adjustments resulting from tax appeals. Households show significant variation in the extent to which they benefit from public education, which is important for our research design. For example, households *with* children accounted for 25.5% of the sample and households *without* children accounted for the remaining 74.5%. We also find significant variation in how the recapture system affects school districts in our sample, with some diverting as much as 57% of their districts'

[4] For a history of property tax recapture in Texas, see for example Villanueva (2018).

[5] In 2021, households in Dallas County could file a protest between April 16, 2021, and May 17, 2021. For more details about how tax protests work, see Section 2.3. For more discussion of the context of tax appeals, see for example Nathan et al. (2020) and Jones (2019).

property taxes and others receiving as much as 23% additional funds from other districts.[6] In terms of protest behavior, for reference, 30.1% of the homeowners in the control group (i.e., those who did not receive any information on school taxes or on recapture) filed a tax appeal in 2021.[7] These tax protests are consequential. For instance, 65.4% of protests lead to a decrease in assessed home value, resulting in average tax savings of $579 in the first year alone.

The results from the first treatment arm indicate that even though the information is publicly available and easily accessible, most households have misperceptions about their respective school shares. When provided with factual information, we observe that households strongly update their beliefs. We leverage the information shocks from the experiment to estimate the causal effects of these beliefs. We find effects that are consistent with predictions of the benefit-based taxation framework. Upon learning that their school shares are higher, households *with* children become *less* likely to protest, whereas households *without* children become *more* likely to protest. The effects of the perceptions about government spending are statistically significant and significant in magnitude. Our baseline estimates imply that increasing the (perceived) school share by 10 percentage points (pp) would cause a drop of 3.67 pp in the probability of filing a protest among households *with* children and an increase of 2.78 pp in the probability of protesting among households *without* children. The effects amount to 11% and 10% of the corresponding baseline protest rates, respectively. These results are robust to a host of alternative specifications and falsification tests.

To assess whether the results were surprising or predictable, we conduct a forecast survey using a sample of 56 experts, most of whom are professors doing research on these topics. After receiving a brief explanation of the experiment, the experts are asked to forecast the experimental findings. The comparison to the expert forecasts indicate that the results were largely surprising, as only a small minority of experts accurately predict the experimental findings.

The second treatment arm explores beliefs about recapture. We find that respondents have significant misperceptions about the recapture share and that they update their beliefs significantly when provided with information in the experiment. Both the levels of misperception and updating, however, are smaller relative to the corresponding findings for the school share, implying that the information shocks for school recapture are not as strong as those for school share. We do not find any statistically significant evidence that perceptions about the recapture share affect the probability of filing a tax appeal, either for households *with* children or *without* children. However, these null effects must be taken with a grain of

[6] The reported numbers refer to net transfers.

[7] Owners can protest directly on their own, which is the main focus of this paper, or they can hire an agent to protest on their behalf. In our control group, 30.1% of owners protest directly and 4.8% use an agent.

salt due to their limited statistical precision.

This paper relates to a recent and growing literature on the role of tax morale (Luttmer and Singhal, 2014). To improve causal identification, this literature uses field experiments consisting of sending messages to taxpayers and then measuring the effects of these messages on tax compliance (Slemrod, 2019).[8] For instance, Blumenthal et al. (2001), Castro and Scartascini (2015), Bott et al. (2020), and Bergolo et al. (2021) each include a treatment arm with a message related to government services. More precisely, these moral suasion messages highlight the importance of taxes to the provision of community services. Evidence on the effects of such messages is mixed.[9] In another related experiment, Carrillo et al. (2021) treated a sample of 400 taxpayers from an Argentine municipality with a joint intervention that recognized them publicly for their good behavior and awarded them with the construction of a sidewalk near their homes. They provide evidence that this intervention decreases subsequent tax delinquency.[10]

We provide three main contributions to the literature. All three contributions are possible due to the novel research design linking data from a survey experiment to administrative tax data at the individual level, which is new to this stream of literature and rare even in broader economics research (Bergolo et al., 2020). The first contribution is to measure and study *perceptions* of government spending. By collecting survey responses from taxpayers, we can precisely measure their beliefs about where their tax dollars go and how those beliefs change in response to accurate feedback.[11]

Our second contribution is the study of the causal mechanisms at play. The experimental interventions in the literature are designed as nudges, combining features that complicate interpretations of the effects. For example, the message in the previously mentioned study by Bergolo et al. (2021) includes normative language describing factual information about government spending, as well as factual information about tax evasion, and the intervention

[8] Beyond field experiments, researchers have studied the role of tax morale on tax compliance using quai-experiments and laboratory experiments. For instance, related evidence from a quasi-experiment (Cullen et al., 2020) and a laboratory experiment (Huet-Vaughn et al., 2019) shows that partisan alignment with the government affects tax compliance.

[9] Blumenthal et al. (2001) find their message does not have a significant effect on tax evasion. Bergolo et al. (2021) and Bott et al. (2020) find significant negative effects in the first year, but the effects do not persist after a year. Castro and Scartascini (2015) find insignificant effects on average, although they provide evidence of heterogeneous effects.

[10] In a related study, Cait et al. (2018) show evidence from a laboratory experiment that tax compliance increases when individuals have the opportunity to voice their preferences for how taxes are to be spent. Also related to our study, Nathan et al. (2020) provides evidence that, conditional on their households' own tax rates, taxpayers' perceptions about the average tax rate affects their decisions to file a protest. Although this result does not pertain to perceptions of government spending, it constitutes consistent evidence that fairness concerns play significant roles in the decision to file a tax appeal.

[11] Beyond tax compliance, recent quasi-experimental evidence demonstrates how the salience of government expenditures can affect electoral outcomes (Huet-Vaughn, 2019; Ajzenman and Durante, 2022).

from Carrillo et al. (2021) awards some taxpayers with the construction of a sidewalk near their homes, as well as a social reward recognizing them publicly for their good behavior. Our experiment, which revolves around factual information on government spending, is designed to test more precisely the mechanism of benefit-based taxation.

Third, we make a methodological contribution. While previous studies measure intention-to-treat effects, our research design allows us to measure treatment effects on the treated. This approach is especially important because, as highlighted by Luttmer and Singhal (2014), the null effects of tax morale messages on behavior may be mechanical due to a null effect on beliefs. For example, consider an intervention consisting of informing taxpayers about their school share. Whether we expect them to react positively, negatively, or not at all depends on whether they underestimate, overestimate, or correctly estimate their school shares. Our research design measures beliefs prior to the information provision, allowing us to fully disentangle reactions based on updating beliefs upwards, downwards, or not at all.[12] Moreover, the effects of this intervention would depend on whether the taxpayers have children enrolled in public schools.[13] We can disentangle those effects because we observe whether the taxpayer has children enrolled in local public schools.[14]

Our study also relates to a small but growing literature on the interplay between tax policy and normative considerations. Benefit-based reasoning was a prominent, and at times leading, approach among tax theorists in the early twentieth century (Seligman, 1908; Musgrave, 1959).[15] However, it has been largely ignored by the modern optimal taxation literature, which instead focuses exclusively on efficiency (Weinzierl, 2018; Scherf and Weinzierl, 2020). Besides benefit-based taxation, a growing body of work is trying to incorporate other normative considerations into tax policy design (Mankiw and Weinzierl, 2010; Weinzierl, 2014; Saez and Stantcheva, 2016; Lockwood and Weinzierl, 2016).[16] This literature is new and mostly theoretical, with empirical evidence limited to survey data, such as asking individuals to choose between hypothetical tax policies (Weinzierl, 2014; Saez and Stantcheva, 2016). We contribute to this literature by providing the first revealed-preference evidence from a

[12] While they do not have survey data to conduct the same analysis, Castro and Scartascini (2015) provides some consistent evidence based on heterogeneity analysis based on the quality of public services across different areas.

[13] Indeed, an additional contribution of our study is that it focuses on public school spending, which is arguably one of the most important government services but receives no attention in this literature. For example, in fiscal year 2019, the U.S. government spent $752.3 billion on its 48 million children in public schools (2019 Public Elementary-Secondary Education Finance Data, U.S. Census Bureau).

[14] The use of survey data is crucial in this regard because this information is difficult to obtain from government records due to data privacy laws regarding children.

[15] Benefit-based taxation differentiates from the typical mechanisms in the literature on the role of tax morale on tax compliance in that the tax revenues fund a public good that benefits the compliant taxpayer directly (Weinzierl, 2018).

[16] For instance, the normative considerations related to equality of opportunity or poverty alleviation.

natural, high-stakes context.

Because protests reveal households' willingness to pay taxes, our study provides causal evidence on the roots of tax compliance and state capacity. Additionally, our results have some direct implications for fiscal and tax policies. First, our results suggest that governments may be able to boost tax compliance by improving perceptions about how taxpayers benefit from their tax dollars, such as through transparency and accountability policies. Second, our evidence suggests that normative considerations, such as benefit-based taxation, should be incorporated into tax policy design. Regardless of its efficiency, a policy that is perceived as unfair by voters and policy makers is unlikely to be feasible in practice. A better understanding of why taxpayers perceive tax policies as fair or unfair can aid in the design of efficient and viable policies.

The rest of the paper proceeds as follows. Section 2 describes the institutional context. Section 3 presents the conceptual framework. Section 4 discusses the experimental design and implementation. Sections 5 and 6 present the results. The last section concludes.

2 Institutional Context

2.1 Property Taxes and Public Schools

In Dallas County, property taxes fund various public services, such as schools, parks, roads, and police and fire departments. In 2021, the average home in Dallas County was worth $327,690, and the average estimated property tax bill was $6,370, implying an effective tax rate of 1.94%. Texas does not have a state income tax. To compensate, revenues from property taxes fund a greater share of local government services in Texas than in many states. School districts receive the largest share of a household's property tax, accounting for nearly half (49.78%) of the average total property tax bill.[17] The second-highest component is the city tax (accounting for roughly 28% of property taxes), followed by hospital (10%), county (8%), college (4%), and special district (<1%) taxes.[18]

Dallas County has 14 main Independent School Districts (ISDs).[19] Homeowners who live within the geographical boundaries of a given ISD jurisdiction are subject to the tax rate for that ISD.[20] Households also have the right to send their children to the K–12 public school(s)

[17] Variation in the school share across households ranges from 13.2% (1st percentile of the distribution) to 90.8% (99th percentile).

[18] See Appendix B.1 for more details.

[19] The total number of ISDs is sixteen, but two of them are extremely small and thus are excluded of the analysis. See Appendix B.2 for more details.

[20] School districts in Texas can set their own tax rates, but they must abide by certain state regulations. See Appendix B.1 for more details.

in their ISD. All households must pay school taxes, regardless of whether they have children enrolled in public schools. The public schools in Dallas County are generally of good quality, although significant differences exist.[21] Alternatively, homeowners can send their children to private schools, conduct homeschooling, or enter a lottery for a chance to send their children to a charter school.[22] Sending children to private schools can be expensive, however. The average tuition cost for private schools in Dallas County is $12,374 per student as of 2022.[23] According to data from the 2020 U.S. Census, about 90% of K–12 students in Dallas County attend a public school.[24]

2.2 Property Tax Recapture

To make public school funding more equitable across school districts, Texas enacted a redistribution system in 1993, called the "Recapture Plan" or "Robin Hood Plan", to divert school tax funds from "property-wealthy" districts to "property-poor" districts.[25] The recapture system has been amended several times since its inception, including a change in 2019 that slowed down the strong growth in the amount recaptured. Nevertheless, redistribution amounts remain substantial under the current recapture formula (Texas Education Agency, 2021c).[26]

In this paper, we focus on the *net* redistribution, which is the difference between the taxes recaptured by the state from wealthy districts and whatever amount is distributed from that state pool to school districts (for specifics on the recapture formula and this calculation, see Appendix B.1). Wide variation in the recapture share occurs across the 14 ISDs that we study. Four ISDs are net givers: the highest giver is Highland Park ISD, which has 57.3% of its school taxes diverted. The remaining ten districts are net receivers: the highest receiver is Mesquite ISD, which receives an additional 23.3% in funding from property taxes diverted from other districts.

[21] For example, according to `www.GreatSchools.org`, 100% of the schools in the Highland Park ISD have above-average ratings in Texas, whereas 43% of schools in the Mesquite ISD have below-average ratings (data accessed on November 4, 2021).

[22] Charter schools are tuition-free public schools that receive funding directly from the state and do not receive funding from property taxes.

[23] Data accessed from https://www.privateschoolreview.com/ exas/dallas-county on January 5, 2022.

[24] More precisely, 89% of kindergarten students and 92.5% of students in grades 1–12.

[25] This system was the result of poor school districts legally challenging the system of state school finances in the late 1980s and early 1990s on state constitutional grounds.

[26] See Appendix B.1 for more details.

2.3 Tax Protests

Each year, the Dallas Central Appraisal District (DCAD) conducts market value appraisals for all homes in the county. Each appraisal results in a "proposed value" for the home, which is an estimate of the home's market value as of January 1 of that year. The DCAD makes this information available to all homeowners through its website and/or by mail.[27] The notice includes additional information, such as the estimated taxes due based on the property's proposed values and how property taxes are allocated across jurisdiction types (e.g., school taxes, city taxes). Homeowners have one month from the date of the notice to protest if they disagree with the proposed assessment value. After the notifications are sent, households have a month from the notification date to file a protest. In 2021, the DCAD presented the proposed values on April 16; as a result, the deadline to protest was May 17.

Homeowners can file a protest by mail using a form included with their mailed notice, or they can file a protest online using a simple tool called uFile.[28] After reviewing the argument, the DCAD can (and often does) inform the homeowner by mail or phone that it will reduce the assessed home value and corresponding tax due. If the homeowner refuses the offer or the DCAD does not offer a settlement, the appeal proceeds to a formal hearing with the Appraisal Review Board.[29] Once protests are resolved, the new tax amount becomes payable either immediately or at some later date (e.g., on October 1st in 2021). Any unpaid taxes eventually become delinquent (e.g., unpaid 2021 property taxes became delinquent on January 31, 2022).

A key feature of this setting is the difficulty in estimating home market values for homes that have not been sold recently, a process involving significant ambiguity and subjectivity. To avoid costly in-person appraisals, the DCAD uses statistical models and large datasets (e.g., recent home sales) to formulate an estimated market value for each property. However, even multibillion-dollar companies like Zillow and Redfin have a hard time estimating market values using statistical models (Parker and Friedman, 2021). This ambiguity in home value is important for the interpretation of our results because it implies that households are not trying to "correct" estimates from the DCAD. Instead, they are presenting a data point (e.g., the sale price of a neighboring home) to support their protest. This distinction is consistent

[27] A sample notification, called the "Notice of Appraised Value", is shown in Appendix G. This notification is available online for every household, and it is also sent by mail to some households (e.g., households with proposed values that increased since the previous year).

[28] To protest online, homeowners need to look up their account (e.g., searching for their own names or addresses) and then follow some straightforward steps in the uFile system. To protest by mail, households who received a notification from the DCAD can use the protest form included with the notification, and households that did not receive a notification can file by mailing a printed form that can be obtained online on either the DCAD's or the Texas Comptroller's website. In 2020, about 75% of direct protests were filed online while the remaining 25% were filed by mail (Nathan et al., 2020).

[29] Homeowners can contest the Appraisal Review Board's decision in court.

with what was expressed in our conversations with officials from some of the county appraisal districts in Texas. Their prevailing view is that households use the subjective nature of the appraisal process as an excuse to complain about their taxes being too high (for more details, see Nathan et al., 2020) and not necessarily to complain about the county's estimate of their home value.

3 Conceptual Framework

To formalize the benefit-based logic, we introduce a simple model of how the provision of government services and redistribution affects the decision to file a protest. Let subscript $j \in \{C, NC\}$ represent the two types of households: those with children enrolled in public schools ($j = C$) and those without ($j = NC$). Let P_j be the outcome of interest: the probability that the household files a tax protest, which is a proxy for its (un)willingness to pay taxes. Let B_j be how much households in group j benefit from each dollar spent in government services. Consider the following relationship:

$$P_j = \gamma \cdot B_j \tag{1}$$

Motivated by the theory of benefit-based taxation, we assume $\gamma < 0$: that is, households benefiting more from government expenditures are less likely to protest their taxes. Let S be the government expenditures in the local public school district and NS be the government expenditures in other local government services (e.g., police, parks, and roads). The two types of households benefit from the two types of government expenditures in the following manner:

$$B_C = \alpha^S \cdot S + \alpha^{NS} \cdot NS \tag{2}$$

$$B_{NC} = \alpha^{NS} \cdot NS \tag{3}$$

where parameters α^S and α^{NS} capture how households benefit from different types of expenditures. The parameter α^S denotes how much a household *with* children enrolled in public school benefits per each dollar spent in public schools. α^{NS} denotes how much households (regardless of whether they have children) benefit per each dollar spent in non-school government expenditures. The key assumption is that households *with* children in public schools benefit more from school expenditures than from non-school expenditures: $\alpha^S > \alpha^{NS}$. This assumption is meant to represent the fact that unlike the benefits from non-school expenditures (e.g., police, roads), which are spread over the entire community, the benefits from school expenditures are concentrated on a subset of the population (households *with* children enrolled in public schools) and thus the members of that subset enjoy them more.

Next, we conduct a simple normalization. Let $G = S + NS$ denote total expenditures and $s = \frac{S}{G}$ denote school expenditures as a fraction of total expenditures, which we previously defined as school share. For the sake of simplicity, we do not incorporate misperceptions into this simple framework. In practice, however, the "s" that matters is the one perceived by the taxpayer when deciding whether to protest. We thus can re-write equations (2) and (3) as follows:

$$B_C = G \cdot (\alpha^S \cdot s + \alpha^{NS} \cdot (1 - s)) \tag{4}$$

$$B_{NC} = G \cdot \alpha^{NS} \cdot (1 - s) \tag{5}$$

Combining equations (1), (4), and (5), we obtain the following:

$$P_C = \gamma \cdot G \cdot (\alpha^S \cdot s + \alpha^{NS} \cdot (1 - s)) \tag{6}$$

$$P_{NC} = \gamma \cdot G \cdot \alpha^{NS} \cdot (1 - s) \tag{7}$$

Using equations (6) and (7), we can see what happens to protest rates if the school share increases:

$$\frac{\partial P_C}{\partial s} = \gamma \cdot G \cdot (\alpha^S - \alpha^{NS}) < 0 \tag{8}$$

$$\frac{\partial P_{NC}}{\partial s} = -\gamma \cdot G \cdot \alpha^{NS} > 0 \tag{9}$$

The intuitions are straightforward. Households *with* children benefit most from school expenditures. Thus, an increase in s implies that they benefit more from government services and that their probability of protesting decreases. In contrast, households *without* children do not benefit from school expenditures. Thus, when s increases, their benefits from government services go down and their probability of protesting goes up. Moreover, if we subtract equation (9) from (8), we obtain the following:

$$\frac{\partial P_C}{\partial s} - \frac{\partial P_{NC}}{\partial s} = \gamma \cdot G \cdot \alpha^S < 0 \tag{10}$$

In other words, the difference in the effect of s between households *with* children versus those *without* children can be tracked to a key parameter of interest, α^S, which is how much households *with* children benefit from school expenditures.

Prediction 1: *An increase in the school share should negatively affect the protest probability of households with children in public schools and positively affect the protest probability of* *households without children in public schools.*

This setup corresponds to the simplest case and is based on two simplifying assumptions. First, it assumes that households are entirely selfish and that households *without* children do not benefit at all from school spending, although in practice these taxpayers may feel good about helping other parents in the community. Second, it assumes that benefits from non-school services are the same for households *with* children as for households *without* children in public schools. We choose this setup due to its simplicity, but in Appendix A, we show that our prediction still holds under more general assumptions.

It is straightforward to extend this simple model to include redistribution of school taxes. For the sake of brevity, we consider the analysis from the perspective of a household in a wealthy school district whose school taxes are redistributed to disadvantaged school districts.[30] Let $r \in [0, 1]$ represent the fraction of school taxes that are transferred from the household's own school district to other school districts, which we previously defined as the recapture share. For instance, $r = 0.4$ would indicate that 40% of school taxes are redistributed to other school districts. We can extend equations (2) and (3) to incorporate recapture into the model:

$$B_C = \alpha^S \cdot S \cdot (1 - r) + \alpha^{NS} \cdot NS \tag{11}$$

$$B_{NC} = \alpha^{NS} \cdot NS \tag{12}$$

We normalize equations (11) and (12) by total expenditures, combine them with equation (1), and then rearrange them as follows:

$$P_C = \gamma \cdot G \cdot (\alpha^S \cdot s \cdot (1 - r) + \alpha^{NS} \cdot (1 - s)) \tag{13}$$

$$P_{NC} = \gamma \cdot G \cdot \alpha^{NS} \cdot (1 - s) \tag{14}$$

Using these equations, we can see what would happen if we increase the recapture share:

$$\frac{\partial P_C}{\partial r} = -\gamma \cdot G \cdot s \cdot \alpha^S > 0 \tag{15}$$

$$\frac{\partial P_{NC}}{\partial r} = 0 \tag{16}$$

The intuitions are straightforward. For households *with* children, more recapture means fewer benefits for their local school district and thus less willingness to pay taxes. Households

[30] The forces at play are similar from the opposite perspective, wherein a disadvantaged district receives funds from more advantaged districts.

without children in the school district do not benefit from school taxes, regardless of which school district receives the funding, so their willingness to pay taxes is unaffected by recapture.

We can also subtract (16) from (15) to show the following:

$$\frac{\partial P_C}{\partial r} - \frac{\partial P_{NC}}{\partial r} = -\gamma \cdot G \cdot s \cdot \alpha^S > 0 \tag{17}$$

Again, the difference in effects between households with children and without children is determined by parameter α^S.

Prediction 2: *An increase in the recapture share should increase the protest probability for households with children in public schools, but it should not affect the protest probability for households without children in public schools.*

This framework assumes that households are totally selfish and care only about how they benefit from government services. In practice, households may appreciate that their tax dollars help the community. In Appendix A, we provide an extension of this framework that incorporates such altruism. Prediction 2 thus must be taken in context because it no longer holds once we allow for altruism.

4 Experimental Design and Implementation

4.1 Subject Recruitment

We mailed our letters so that they would be delivered close to the time that homeowners in Dallas County could start filing tax appeals. Appendix C shows a sample envelope, and Appendix D shows a sample letter. We included several features to indicate the legitimacy of the letters. For example, the letters were sent on behalf of researchers at The University of Texas at Dallas, a well-known institution in Dallas County. The envelope featured the school's logo, the name of a professor from that university, and non-profit organization postage. The letter itself included a physical address for the researcher and a link to the study's website (see Appendix E for a screenshot of the website). It also provided contact information for the researchers and institutional review board. The letter salutation included each recipient's name, and recipients' names and addresses were printed at the bottom of the second page so that they appeared through the envelope window. In cases where properties were jointly owned by multiple individuals (typically, husband and wife), we sent one letter to the address but listed all owners on the letter. As previously mentioned, the letter also mentioned the recipient household's proposed value and estimated property tax amount for 2021.

Most importantly, our letters included an invitation to participate in an online survey and included the URL of the survey. Each subject was asked to enter a unique survey code, which was included in the letter right next to the survey URL. This code allowed us to identify survey respondents and link their responses to the administrative records. In addition to the opportunity to contribute to a research study, we included two additional incentives for survey participation. First, the letters indicated that detailed, step-by-step instructions on how to file a protest online or by mail would be provided at the end of the survey.[31] As a second incentive, subjects were informed that survey respondents would be entered into a raffle for 20 prizes worth $100 each.[32]

4.2 Survey Design

In this section, we summarize the main features of the survey.[33] We start by asking a critical question, that is, whether the respondent's household has children enrolled in grades K–12 at their local public school district, and if so, how many. This critical information is missing from administrative records of the tax agency and thus the analysis would be impossible without this question, particularly the heterogeneity analysis concerning the benefit-based taxation framework, which is the main form of heterogeneity that we anticipate in the RCT pre-registration.

The module about school taxes can be summarized as follows:

- **Step 1 (Elicit Prior Belief):** We begin by providing the estimated total property tax amount of the respondent's home in 2020 (based on administrative records). We then explain that this total amount is the sum of different components, such as school, city, and hospital taxes. We ask respondents to guess their school share in 2020, using any amount between 0% and 100%.

- **Step 2 (Information-Provision Experiment):** For every subject, we calculate the "correct" answer to the previous question based on administrative records. We then randomize whether the subject sees the correct answer. Each subject faces a 50% probability of being shown this information. To avoid respondents making inferences from the act of receiving information, we make the randomization explicit. On the

[31] This walkthrough included hyperlinks to relevant websites and screenshots of a sample protest using information for a fictitious household for added clarity. To access these instructions, subjects were provided with a URL and a code on the final screen of the survey. A copy of the web instructions is included in Appendix F. Nathan et al. (2020) show that these instructions have a significant positive effect on the probability of protesting.

[32] All respondents were entered into the same raffle, but only a random half of respondents were informed about the raffle in the letter (i.e., before deciding whether to participate in the survey). This randomization aimed to assess the effectiveness of raffle prizes in increasing response rates, which can be useful information for future researchers conducting similar field experiments. For more details, see Appendix B.5.

[33] A sample of the full survey instrument is attached as Appendix H.

first screen, we inform respondents that some participants will be randomly chosen to receive the information and that they will find out on the next screen if they are selected. On the following screen, we inform subjects whether they are chosen to receive the feedback.

- **Step 3 (Elicit Posterior Belief):** We give all subjects the opportunity to revise the guess they provided in Step 1. To avoid asking the exact same question about their 2020 taxes (i.e., the year prior to our intervention), we instead ask about their 2021 taxes (i.e., the most recent year). To avoid subjects making inferences based on the opportunity to re-elicit their guesses (e.g., subjects inferring that we ask again only if their answer in step 1 is incorrect), we explicitly inform them that all survey participants have this opportunity, regardless of their initial guesses.

To learn about the causal effects of beliefs, it is critical to leverage information on prior beliefs. When provided with feedback during the information-provision experiment, individuals who underestimate may update their beliefs upward and those who overestimate may adjust their beliefs downward. Some individuals may have accurate priors and thus may not make any updates. Whether an individual's probability of protesting increases, decreases, or remains the same should depend on the individual's beliefs before receiving the information. For this reason, we conduct the information-provision experiment within the survey, as opposed to providing the information directly in the letter, to measure beliefs prior to information provision. To leverage the effect of the information on prior beliefs, we use the same econometric models used in other information-provision experiments (see e.g., Cullen and Perez-Truglia, 2022; Bottan and Perez-Truglia, 2020b).

The following module is about the recapture share.[34] Some subjects may not know about or understand recapture. Thus, we start with a couple of short paragraphs summarizing the recapture system. The rest of the module follows the same structure as previously described for steps 1 through 3. We elicit beliefs about the recapture share in two steps. First, we ask respondents to guess if their school district will receive more, the same, or less in taxes than what households in their district paid in school taxes. The second step is quantitative in nature. If the respondent selects "More" (or "Less") in the first question, we ask them to guess how much *more* (or *less*) funding their school district will receive as a share of the district's school tax revenues due to recapture, using any amount between 0% and 100%. We then conduct step 2 (information-provision experiment) and step 3 (elicitation of posterior beliefs).

[34] Note that the recapture share is ISD-specific, whereas the school share is household-specific.

We cross-randomize subjects to receive the two pieces of information about school taxes and recapture, respectively, with a 50% probability for each. Thus, roughly 25% of the sample receives both pieces of information, 25% receives the first piece of information only, 25% receives the second piece of information only, and 25% receives no information at all.

These questions comprise the core of the survey. We also include a series of additional questions, including one question that serves as a secondary outcome in the analysis of the effects of beliefs. We ask respondents whether they plan to file a protest this year in a 1-4 likelihood scale. This outcome allows us to pick up short-term effects on the *intention* to protest, even if those effects do not materialize into actual protests. For descriptive purposes, we include questions asking respondents' gender, age, ethnicity, education, and political party. To provide supplemental evidence, towards the end of the survey, we include additional questions that are described in more detail in the following sections.

4.3 Subject Pool

We sent the previously described letters to 78,128 households representing a subsample of the universe of all households in Dallas County, Texas. We arrived at this subsample by applying several filters (e.g., excluding commercial properties and non-owner-occupied residences.)[35] When selecting this sample, we stratified the randomization at the ISD level to ensure wide representation of the beneficiaries and contributors of the recapture system.[36] We can link each survey respondent to rich sources of administrative data, including whether the subject protested in any year from 2016 to 2020, as well as detailed information on property ownership, address, number of bedrooms and other features, exemption amounts, taxable values, and tax rates.

We timed the intervention so that our letters would arrive early enough before the protest deadline to influence the recipient's decision. We created the letters on April 16th, 2021, as soon as the administrative data, including 2021 proposed values, became available. To accelerate delivery, we used a mailing company in Dallas County (i.e., the same county as all recipients). The mailing company dropped off the letters at the local post office on April 20, 2021, and estimated that most would be delivered in the next couple of days. Consistent with this projection, we began to receive survey responses and visits to the study's website on April 22, 2021.[37]

Of 78,128 households invited to the survey, 2,966 answered the first two questions and

[35] For the full inclusion criteria, see Appendix B.2.
[36] For more details, see Appendix B.3.
[37] More details about the timing of survey responses are provided in Appendix B.6.

2,821 completed the two key modules on posterior beliefs about the recapture share.[38] The implied response rate of 3.6% ($= \frac{2{,}821}{78{,}128}$) is comparable to the response rate of 3.7% from a previous study in this same context and with a similar recruitment method (Nathan et al., 2020). Moreover, the response rate of 3.6% is on the same order of magnitude as the response rate of surveys that use this recruitment method (4.7%, as reported in Sinclair et al., 2012).[39] Among respondents, the median time to complete the survey was 11.3 minutes. Towards the end of the survey, we included an attention check similar to the one used in other studies (Bottan and Perez-Truglia, 2020a), which 92.1% of respondents successfully passed. This passing rate is high in the context of survey studies, particularly given that the attention check was located at the very end of the survey when fatigue was likely at its highest.

Of the 2,821 survey responses, we drop responses that, as explained in the RCT pre-registration, could not be excluded ex ante because of data availability. We drop 36 responses from subjects who, according to the DCAD's records, had already filed a protest before starting our survey and 23 additional subjects who responded to the survey after the deadline to file a protest, as the survey information could not have affected their decisions to protest. We similarly drop 185 subjects who, according the DCAD's records, had already hired a tax agent before starting our survey.[40]

When studying perceptions via survey data, it is important to deal properly with outlier beliefs. Some individuals may provide guesses that are wildly inaccurate not because they truly hold such extreme beliefs but because they misunderstand the question, make a typo, or just do not pay attention to the question. The "information shocks" for these individuals can be large but meaningless, which can create a significant attenuation bias. To reduce sensitivity to outliers, we follow the standard practice in information-provision experiments and drop respondents with the most extreme misperceptions in their prior beliefs (see e.g., Fuster et al., 2021; Cullen and Perez-Truglia, 2022; Bottan and Perez-Truglia, 2020a). For the baseline specification, we use a conservative definition of outliers that drops 467 subjects from the bottom 5% and top 5% of the distribution of prior misperceptions.[41] After applying these filters, 2,110 respondents remain, constituting our main subject pool. Since these exclusions are based on pre-treatment variables (e.g., prior beliefs), they should not compromise the validity of the experimental variation. As a robustness check, we reproduce the analysis with more lax definitions of outliers (results presented in Section 5.5). Finally, we provide several sharp falsification tests to address any potential concerns about the internal validity of the

[38] See Appendix B.3 for more details about the sample and Appendix B.6 for more details about attrition rates and balance tests.

[39] The 4.7% response rate corresponds to a mailing of a personally addressed postcard inviting a household to complete a web-based survey using a unique alphanumeric code.

[40] For more details, see Appendix B.2.

[41] For more details on the distribution of outlier observations, see Appendix B.7.

results, such as event-study analyses.

Column (1) of Table 1 presents descriptive statistics about the subject pool. Prior to any adjustment resulting from protests, the average subject owns a home with an assessed market value of $349,988 and property taxes of $7,738 (corresponding to an average tax rate of 2.21%). Around 25.5% of respondents have children enrolled in a local public school, 42.9% are women, 44.3% self-identify as White, 38.3% have a college degree, and on average they are 49.6 years old.

In terms of observable characteristics (e.g., home value, number of bedrooms, or tax rate), the subject pool is similar to the universe of households in the county. Differences between survey respondents and non-respondents are statistically significant but small (see Appendix B.3). However, one significant difference is that, relative to the universe of households, respondents to the survey are substantially more likely to file a protest in 2021 and in previous years. By design, our study targets individuals who would seriously consider protesting, which increases statistical power by securing more variation in the outcome variable.[42] Moreover, our letter describes tax protests, so subjects considering filing a protest in 2021 are likely to pay attention to the letter and thus also likely to notice the survey link included in the letter.[43]

Columns (2) through (5) of Table 1 break down the average characteristics in each of the four treatment groups. All characteristics shown in Table 1 are determined pre-treatment and thus should not be affected by the treatment assignment.[44] Column (6) reports p-values for the null hypothesis that the average characteristics are equal across the four treatment groups. Table 1 shows that, consistent with successful random assignment, the observable characteristics are balanced across treatment groups.[45] In Appendix B.6, we present alternative versions of the randomization balance tests, such as breaking the sample down by households *with* and *without* children. We also show that participation in the survey and attrition among participants is orthogonal to treatment assignment, which is expected given

[42] Specifically, when selecting households to participate in the survey, we over-sample those most likely to protest, such as households with a history of increased estimated taxes. For more details, see Appendix B.2.

[43] Indeed, this higher propensity to protest among survey respondents is consistent with results from Nathan et al. (2020), who use a similar recruiting method to collect survey responses in this same context. Moreover, our letter promises instructions on how to file a protest as a reward for participation, so it is natural that interested respondents would be more likely to participate. Additionally, these instructions likely make it easier for survey respondents to file an appeal, as documented in Nathan et al. (2020).

[44] Some questions, such as gender of the respondent, are asked after the information-provision stage. However, treatment assignment should not affect these responses. For instance, we do not expect information on school spending to change responses regarding gender or education level.

[45] The difference is statistically significant for one of the variables (owner protest in 2020). Given the large number of tests conducted, a few differences may be statistically significant just by chance. To be safe and to follow best practices in field experiments (Athey and Imbens, 2017), we include this variable in the set of control variables in all regressions.

that subjects randomly receive a treatment after they start the survey.

4.4 Outcomes of Interest

As stated in the RCT pre-registration, the main outcome of interest is a dummy variable indicating whether the household protested on its own in 2021.[46] To get a sense of the baseline protest rate, we consider subjects in the control group (i.e., those who do not receive any information on school taxes nor recapture). Approximately 30.1% of those owners file a tax appeal in 2021. These tax protests are consequential: 65.4% lead to a decrease in the assessed home value, among which the average tax savings were $579 in the first year alone.[47]

Owners can file their own protests, which is the main focus of this paper, or they can hire an agent to file a protest on their behalf. In addition to the 30.1% of owners who protest directly, 4.8% use an agent.[48] Due to the nature of the setting and the timing of the protest process, and as stated in the RCT pre-registration, we expect our information to primarily affect whether households file their own protests. Because we provide information to the households and not to their agents, our experiment should not affect the agent's behavior. Moreover, it is unlikely that the information provided in our survey would influence whether households hire an agent. According to anecdotal accounts with households, tax agents and representative from assessor's offices, households typically sign contracts with agents well in advance of the date when the proposed values are announced, which is when we provide the information. Some agents even offer long-term contracts to file protests on behalf of the owner over many years. For these reasons, we study protests through tax agents as a separate outcome, in a falsification spirit. For the sake of brevity, in the rest of the paper we use the term "protest" as shorthand for direct protests by the homeowner, unless explicitly stated otherwise.

4.5 Expert Prediction Survey

To assess whether the experimental results are surprising, we conduct a forecast survey with a sample of experts. A sample of the full survey instrument is attached as Appendix I. In this survey, which follows best practices (DellaVigna et al., 2019), we describe the experiment and ask experts to forecast the key results in a way that is comparable to the experimental

[46] The protest variable is based on data downloaded from the DCAD website on June 22, 2021.

[47] These calculations are based on data downloaded from the DCAD website on December 2021. The remaining protests are either unresolved by December 2021 (12.2%) or resolved with no change in the assessed home value (22.4%).

[48] These statistics from 2021 refer to the control group: those who did not receive any information on school taxes or on recapture.

estimates. More precisely, we elicit their prediction of the effect of a 10 pp shock to the belief about the school share, separately for households *with* and *without* children. We then conduct the corresponding elicitations for beliefs about the recapture share.

We collected responses from experts in two ways. First, we posted the survey on the Social Science Prediction Platform from July 13, 2021, to December 31, 2021. Second, on November 2021, we emailed an invitation to the prediction survey directly to a list of 238 professors with publications related to our experiment. The final sample includes 56 experts. Of these, 21.4% responded to the survey through the Social Science Prediction Platform, and the remaining 78.6% responded through our email invitation.[49] The final sample thus comprises 82.1% professors, 12.5% PhD students, 3.6% post-docs, and 1.8% researchers. Most (78.6%) are from the field of economics; 66.1% report having done research on taxation and 25% on preferences for redistribution.

5 Perceptions about School Spending

5.1 Accuracy of Households' Prior Beliefs

Transparency and accountability efforts have made information about property taxes publicly available. Each year in Dallas County, the Dallas Central Appraisal District (DCAD) provides homeowners with a Notice of Appraised Value, which contains a detailed break-down of the household's property taxes by tax jurisdiction, including the share of their property taxes that funds public schools.[50] We acknowledge that the ease of access to this information does not mean that everyone searches for it or uses it. Many other contexts show that individuals often misperceive easily accessible information, such as the official inflation rate (Cavallo et al., 2017) or recent trends in national home prices (Bottan and Perez-Truglia, 2020a).

Figure 1(a) shows a histogram of the degree of misperceptions about the school share.[51] The x-axis corresponds to the difference between the actual school share (i.e., potential feedback) versus that perceived by respondents. For the sake of brevity, we use the term feedback

[49] Among the responses from the Social Science Prediction Platform, we exclude respondents who are not academics, who do not have a PhD, or who are not pursuing a PhD.

[50] See Appendix G for a sample of this notice, with the breakdown by tax jurisdiction shown on the second page. The county uses the prior year's jurisdictional tax rates to estimate taxes due in the Notice of Appraised Value because the tax rates for the current year are set later in the year. In practice, tax rate changes are uncommon, so approximation errors are typically negligible. In our study, we use the same definition of estimated taxes because these are the relevant object of study and they represent the subjects' best approximation at the time of deciding whether to protest.

[51] All results are based on the final survey sample, which excludes the outlier misperceptions (i.e., the bottom and top 5%). Including the extreme observations would increase the degree of misperceptions; for more details, see Appendix B.7.

to refer to potential feedback. A minority of subjects have accurate perceptions: more precisely, 32.6% of subjects guess the school share to be within ± 5 pp of the actual school share. Misperceptions are quite large on average: the mean absolute error is 16.57 pp. The large degree of misperceptions implies sufficient scope for the information provision experiment to shock beliefs. Another interesting feature of prior beliefs is that the misperceptions show a systematic bias: on average, subjects underestimate the school share by 13.08 pp, as indicated by the mean error. This systematic bias is quite noticeable in Figure 1(a), where more observations fall on the the right half of the histogram (corresponding to underestimation) than on the left half (corresponding to overestimation).

5.2 Belief Updating

We find that taxpayers update their inaccurate beliefs when provided with accurate feedback. To model belief updating, we use a simple Bayesian model that has been shown to accurately represent belief formation in other information-provision experiments on a wide range of topics, such as inflation expectations (Cavallo et al., 2017), salary expectations (Cullen and Perez-Truglia, 2022), and home price expectations (Fuster et al., 2021).

We use the subscript i to index the subjects. We use the variable s_i^{prior} to represent subject i's belief right before the information-provision experiment. We use the variable s_i^{feed} to represent the value of the feedback that the subject can potentially receive in the experiment. We define the variable T_i^S as a binary variable that equals 1 if subject i is selected to receive the information and 0 if not. We define variable s_i^{post} as the posterior belief. Specifically, s_i^{post} represents the perceived school share after the taxpayer sees, or does not see, the feedback.

An individual will form her posterior belief (s_i^{post}) as the average of the prior belief (s_i^{prior}) and the feedback (s_i^{feed}), weighted by a parameter α that captures the degree of learning. This parameter can range between 0 (individuals ignore the feedback) and 1 (individuals fully adjust to the feedback), and it is a function of the relative precision of the prior belief versus that of the feedback.[52] This Bayesian updating model can be summarized by the following linear relationship:

$$s_i^{post} - s_i^{prior} = \alpha \cdot \left(s_i^{feed} - s_i^{prior} \right) \tag{18}$$

Intuitively, Bayesian learning predicts that, when shown feedback, respondents who overestimate the school share would revise their beliefs downward, whereas respondents who

[52] These results assume normal distribution of priors and feedback and assume that the variance of the prior and the variance of the feedback are independent of the mean of the prior. For more details, see Hoff (2009).

underestimate the school share would revise their beliefs upward. Figure 1(b) estimates this Bayesian learning model using a binned scatterplot. The x-axis corresponds to the gaps in prior beliefs ($s_i^{feed} - s_i^{prior}$), and the y-axis to the belief updating ($s_i^{post} - s_i^{prior}$). Intuitively, the x-axis shows the maximum revision we would expect if the respondent were to fully react to the information, and the y-axis shows the actual revision. In the case of no updating, the observations should form a horizontal line; in the other extreme, under full updating, the observations should form a 45-degree line. The red circles in Figure 1(b) correspond with subjects who are shown feedback about the school share. Consistent with significant updating, there is a strong relationship between the updated beliefs and prior gaps: an additional percentage point (pp) in perception gap is associated with an actual revision that is 0.809 pp higher.

The gray squares in Figure 1(b) correspond with the subjects who do not receive information about the school share. In the absence of feedback, these subjects should not update their beliefs. However, in practice, individuals might revise their beliefs in the direction of the feedback for spurious reasons even when they receive no feedback. For instance, respondents may reassess their answers or correct typos when asked a question a second time, leading to an answer that is closer to the truth. The gray squares indicate a weak relationship between belief updating and prior gaps in the group that was not shown the feedback: an additional 1 pp in the prior gap is associated with an actual revision that is 0.052 pp higher. This effect is statistically significant (p-value $<$0.001) but economically very small. This result is consistent with other information-provision experiments that show evidence of spurious revisions (e.g., Fuster et al., 2021; Cullen and Perez-Truglia, 2022).

We can exploit the random assignment from the information-provision experiment to control for spurious learning:

$$s_i^{post} - s_i^{prior} = \tau + \alpha \cdot \left(s_i^{feed} - s_i^{prior}\right) \cdot T_i^S + \beta \cdot \left(s_i^{feed} - s_i^{prior}\right) + \epsilon_i \tag{19}$$

In this model, parameter α represents true learning arising from the information provision (not spurious learning), whereas parameter β captures spurious learning. Parameter α can be computed from the estimates in Figure 1(b). Specifically, the α parameter corresponds to the difference in the regression slopes between the subjects who are and are not shown the feedback. Since α captures the effect of the exogenous shocks induced by the information-provision experiment, it can be used as an excluded instrument in the econometric model explained in Section 5.3. The estimated α is large ($0.757 = 0.809 - 0.052$) and highly statistically significant (p-value$<$0.001). This difference suggests that a 1 pp information shock induces a 0.757 pp effect in the subject's posterior belief. This shows that, although subjects did not update fully to the feedback, they were close to it. This finding of imperfect updating is consistent with other information-provision experiments and it is likely due to some subjects mistrusting the source of the feedback or simply not paying enough attention to the survey.

Appendix B.6 provides some additional results and robustness checks. For instance, this appendix shows that learning from the feedback is compartmentalized (i.e., subjects do not use the information about school share to update beliefs about the recapture share). This appendix also shows that the belief updating results are similar for households *with* and *without* children.

5.3 2SLS Model

Let P_i^{2021} denote the main outcome of interest: an indicator variable that equals 100 for individuals filing a protest in 2021 (i.e., post-treatment) and 0 otherwise. As discussed in the conceptual model in Section 3, and as noted in the RCT pre-registration, the effects of the school share information treatment on protests are expected to have different signs depending on whether the household has children enrolled in public schools. Let $C_i \in 0, 1$ be an indicator variable that equals 1 if the household has a child enrolled in a local public school and 0 otherwise. Hence, we can use the following econometric specification to estimate our parameters of interest:

$$P_i^{2021} = \beta_0 + \beta_C^S \cdot C_i \cdot s_i^{post} + \beta_{NC}^S \cdot (1 - C_i) \cdot s_i^{post} + \beta_1 \cdot C_i + \epsilon_i \tag{20}$$

The term ϵ_i represents the error. The two parameters of interest are β_C^S and β_{NC}^S. The conceptual model from Section 3 predicts that $\beta_C^S < 0$ and $\beta_{NC}^S > 0$. Moreover, the difference between these two parameters, $\beta_C^S - \beta_{NC}^S$, is of special interest because it captures a key parameter: how much households *with* children benefit from school expenditures (see equation (10) for the case of school share and equation (17) for the case of recapture share). As posterior beliefs (s_i^{post}) are endogenous and thus could suffer from a host of omitted variable biases, we estimate equation (20) using 2SLS, exploiting the exogenous variation in posterior beliefs induced by the information-provision experiment. More precisely, we estimate the following model:

$$\begin{aligned} P_i^{2021} = &\beta_0 + \beta_C^S \cdot C_i \cdot s_i^{post} + \beta_{NC}^S \cdot (1 - C_i) \cdot s_i^{post} + \beta_1 \cdot C_i + \\ &+ \beta_2 \cdot C_i \cdot (s_i^{feed} - s_i^{prior}) + \beta_3 \cdot (1 - C_i) \cdot (s_i^{feed} - s_i^{prior}) + X_i \beta_X + \epsilon_i \end{aligned} \tag{21}$$

The endogenous variables are $C_i \cdot s_i^{post}$ and $(1 - C_i) \cdot s_i^{post}$, for which we use the excluded instruments $C_i \cdot T_i^S \cdot (s_i^{feed} - s_i^{prior})$ and $(1 - C_i) \cdot T_i^S \cdot (s_i^{feed} - s_i^{prior})$.[53]

We can demonstrate the intuition behind the model using a simple example. Consider a pair of subjects with children enrolled in public schools that share the same bias about the school share: both underestimate the actual school share by 20 pp. Suppose we randomly assign information about the true school share to one of them. We expect that, relative to the subject who does not get the information, the subject who receives the information adjusts their perceived school share upward. For the sake of argument, assume that the subject who does not receive the information continues to underestimate the actual school share by 20 pp and that the subject who does receive the information reacts to it by underestimating the school share by just 10 pp. The information provision is thus equivalent to a +10 pp shock to the perceived school share. We can then check the behavior of this pair of households in the weeks after they receive the information. For example, the +10 pp shock to the perceived school share could translate to a lower probability of filing a protest. Again, for the sake of the argument, assume that the +10 pp shock to the belief causes a 2 pp drop in the probability of protesting. Combining these two results, we can estimate that $\beta_C^S = -0.2$, that is, each 1 pp increase in the perceived school share lowers the probability of protesting by 0.2 pp.[54]

The term X_i in equation (21) corresponds to a set of additional control variables. In principle, the 2SLS model leverages the experimental variation, so there is no need to include additional control variables for identification. However, the inclusion of additional control variables can be helpful, for instance, to reduce the variance of the error term and thus improve precision (McKenzie, 2012). The vector of control variables includes basic pre-treatment information, such as the household's prior history of tax appeals.[55]

Following the regression specification we use to study the effects of the school share (equation (20)), it is straightforward to define the regression specification to study the effects

[53] Note that equation (21) controls for the prior gaps in beliefs ($C_i \cdot (s_i^{feed} - s_i^{prior})$) and ($(1-C_i) \cdot (s_i^{feed} - s_i^{prior})$). The inclusion of these control variables ensure that the excluded instruments isolate the information shocks that are driven purely by the random assignment of the feedback (T_i^S).

[54] Typically in 2SLS models, if treatment effects are heterogeneous, the estimates identify the local average treatment effects of beliefs (Imbens and Angrist, 1994). More precisely, in our study, our estimates would give a higher weight to subjects whose beliefs are more affected by the information-provision experiment. By construction, this weight will be higher for subjects with larger prior misperceptions and, conditional on the misperceptions, those who react more strongly to feedback.

[55] The full set of additional control variables includes the log of total market value in 2021, the growth in total market value between 2021 and 2020, an indicator for positive growth, an indicator of whether the property value was re-evaluated in 2021, the 2021 estimated property taxes (in logs), a dummy for homestead exemption in 2021, an indicator for homestead binding in 2021, the household's effective tax rate, a dummy variable for multiple owners, a dummy variable for condos, the total living area, the number of bedrooms, the number of full baths, the building age, a set of dummies for school districts, the survey start date, and indicator variables for whether the household protested in each pre-treatment period since 2016 (one set for direct protests and another set for protests through agents).

of recapture. Indeed, as these two information treatments are cross-randomized for the same sample, we estimate all effects simultaneously in a single 2SLS regression. See Section 6.3 for a discussion of the recapture share estimates.

5.4 2SLS Estimates

The 2SLS estimates for school share are presented in the top half of Table 2.[56] In column (1), the dependent variable is the main outcome of interest: an indicator variable that equals 100 if the subject protests directly in 2021 and 0 otherwise. According to the hypothesis of benefit-based taxation, an increase in the perceived school share should decrease the probability of protesting for households with children (because they find out that they benefit from government services more than they thought), whereas this information should have the opposite effect on households without children. The results are consistent with this hypothesis. The coefficient for households *with* children is negative (-0.367) and statistically significant (p-value=0.096). The coefficient for households *without* children is positive (0.277) and statistically significant (p-value=0.032). Most importantly, the difference between the two coefficients (-0.367 and 0.277) is statistically significant (p-value=0.012).

As a thought experiment, consider what would happen if the perceived school share increases by 10 pp. The estimates from column (1) of Table 2 indicate that this change would cause a drop of 3.67 pp ($= 0.367 \cdot 10$) in the probability of filing a protest for households *with* children and an increase of 2.77 pp ($= 0.277 \cdot 10$) in the probability of protesting for households *without* children. These effects would be roughly equivalent to 11% and 10% of the baseline protest rates (33.86 pp and 28.83 pp, reported at the bottom rows of Table 2).

Column (2) of Table 2 is identical to column (1), except that it uses a different dependent variable: an indicator variable that equals 100 if at the end of the survey the subject responds "very likely" to the question on the likelihood to protest in 2021 and 0 otherwise. This outcome measures the intention to protest and allows us to measure if the effects of the information lead to an intention to protest immediately after the information is provided. For reference, at the time of the survey, 45.4% report that they are very likely to protest (this corresponds to the baseline rate, combining subjects *with* and *without* children who do not receive any feedback). It is important to note that the stated intention to protest correlates significantly with whether the individual actually files a protest, but that correlation (0.410) is far from perfect (the correlation coefficient is 0.410 for the no-feedback group,

[56] We present the 2SLS estimates directly because they can be interpreted more easily. Nevertheless, due to the strong first stage (i.e., high belief updating), the 2SLS estimates are similar (in terms of magnitude and statistical significance) to the reduced-form estimates. For more details, see Appendix B.8.

p-value<0.001).[57] Due to this imperfect correlation, the effects on the intention to protest at the time of answering the survey should not expected to be "mechanically" the same as the effects on actual protests.

The results from column (2) of Table 2 are consistent with the results from column (1). In column (2), the coefficient for households with children is negative (-0.408) and similar in magnitude to the corresponding coefficient from column (1) and statistically significant (p-value=0.080). The coefficient for households *without* children is positive (0.269), on the same order of magnitude as the coefficient from column (1), and statistically significant (p-value=0.062). Again, most importantly, the difference between the coefficients for households with children versus those without children (-0.408 and 0.269) is statistically significant (p-value=0.014).

5.5 Robustness Checks

To probe the robustness of the school share results, columns (3) and (4) of Table 2 provide two falsification tests. The first falsification test uses a dependent variable to indicates whether the household protests through an agent. As explained in Section 4.4, it is highly unlikely that the information provided in our survey would affect protests through an agent. The results are reported in column (3), which estimates the same regression from column (1) but using protests conducted by agents as the dependent variable. As expected, the coefficients from column (3) are close to zero (-0.028 and -0.033) for both households *with* and *without* children), precisely estimated, and statistically insignificant (p-values of 0.816 and 0.518). Moreover, the difference between the coefficients for households *with* and *without* children is close to zero (0.006), precisely estimated, and statistically insignificant (p-value=0.966).

For the second falsification test, we exploit the timing of the information intervention in an event-study fashion. In column (4) of Table 2, we estimate the same baseline regression from column (1), except that it uses as a dependent variable the protest decision in a pre-treatment year (2020), rather than in the post-treatment year (2021). Intuitively, the information-provision experiment should not affect behavior in this pre-treatment year because the individuals are not yet exposed to the information. We thus expect the coefficients from this falsification exercise to be close to zero and statistically insignificant. The results from column (4) confirm our expectations: the coefficients from column (4) are close to zero (0.110 and -0.065, for both households *with* and *without* children), precisely estimated, and statistically insignificant (p-values of 0.545 and 0.504); moreover, the difference between

[57] This correlation (0.410) is computed among subjects who do not receive any feedback. Among respondents who report being very likely to protest, 56.8% end up protesting directly or through an agent. On the other hand, among the individuals who do not report being very likely to protest, 16.8% end up protesting.

households *with* children and *without* children is also close to zero (0.175) and statistically insignificant (p-value=0.398).

We can extend this same falsification test to other pre-treatment years for which we have readily available data. For ease of exposition, the results are presented in graphical form in Figure 2(a). The x-axis denotes the year of the dependent variable (i.e., whether the owner protests directly in years 2016 through 2021). This figure focuses on the main result, corresponding to the difference between households *with* children versus *without* children. As expected, this difference is close to zero and statistically insignificant for each pre-treatment year (2016–2020) and then negative and statistically significant in the post-treatment year (2021).

One usual concern with 2SLS estimation concerns weak instruments (Stock et al., 2002). Given the strong belief updating documented in Section 5.2, weak instruments should not be a concern in our setting. Nevertheless, for a more rigorous assessment, Table 2 reports the Cragg-Donald F statistic, which is commonly used to diagnose weak instruments. The value of this statistic in each regression is well above the rule of thumb of $F > 10$ proposed by Stock et al. (2002): it equals 30.10, 30.22, 30.10, and 30.02, respectively, in columns (1)–(4) of Table 2.[58]

The 2SLS model used for the results in equation (21) assumes a linear relationship between school share and the probability of protesting, a natural starting point due to its simplicity and because it is common practice in the literature on information-provision experiments. To probe that assumption, Figure 2(b) presents a binned scatterplot representation of the reduced-form effects of the information-provision experiment (i.e., not accounting for how the information provision affects prior beliefs). The x-axis corresponds to the interaction between the potential information disclosure and the prior gap (i.e., the excluded instrument). The y-axis corresponds to the probability of protesting in 2021. This binned scatterplot includes all the same control variables used in the 2SLS model. Figure 2(b) seeks to assess whether the relationship between the interaction term on the horizontal axis and the protest probability on the vertical axis is linear, and the figure shows that a linear fit is a reasonable functional form assumption for this context. Additionally, this figure shows that the previously discussed regression results are not driven by outliers.

Table 3 presents additional robustness checks. Columns (1) and (2) of Table 3 reproduce the baseline specification given by columns (1) and (2) of Table 2 for reference. Columns (3) through (10) of Table 3 present the results under alternative specifications. The specification from columns (3) and (4) is identical to the specification from columns (1) and (2), except that we include some additional control variables: the respondent's age, a dummy for individuals

[58] For the results of the first-stage regressions, see Appendix B.8.

that self-identify as White, a dummy for gender, a dummy for college degree, and a dummy for political party (which equals 1 for individuals who self-identify as Democrat).[59] The results from columns (3) and (4) are similar to the baseline results from columns (1) and (2). If anything, these coefficients are slightly larger (-0.714 vs. -0.644 and -0.744 vs. -0.678).

In columns (5) through (8) of Table 3, we try alternative definitions of outliers in prior misperceptions. The baseline specification is already conservative in that it excludes the extreme top and bottom 5% of the distributions. In columns (5) and (6), we use a less stringent definition of outliers based on the upper and bottom 2.5% instead of 5%. The results from columns (5) and (6) are similar to those from the baseline specification of columns (1) and (2), although the coefficients are slightly smaller. In columns (7) and (8), we consider an even more lax definition of outliers, excluding only the upper and bottom 1% of misperceptions. The coefficients from columns (7) and (8) remain consistent with those from the baseline specification of columns (1) and (2), although again the magnitudes are somewhat smaller. These results are consistent with the arguments in Section 4.3 that we should be cautious when including extreme misperceptions because they probably reflect a lack of attention or mistakes, rather than legitimate misperceptions. To explore this further, columns (9) and (10) are identical to the baseline specification from columns (1) and (2), except they exclude respondents who do not pass the attention check included at the end of the survey. Consistent with the attention argument, when we focus on subjects who pass the attention check, the coefficients increase somewhat.

5.6 Comparison to Expert Predictions

Next, we compare our experimental results to expert predictions, as shown in Figure 3. Panel (a) presents the predictions of experts for households *with* children, and panel (b) presents the predictions for households *without* children. The histograms correspond with the distribution of expert predictions estimating the effect of a 1 pp increase in the school share.[60] The solid vertical red line in each panel represents the corresponding estimate from the baseline 2SLS model (column (1) of Table 2), and the red shading denotes the corresponding confidence intervals.

Figure 3 shows that our experimental findings are not obvious to the sample of experts. Our experimental results are consistent with experts who predicted that the school share

[59] These variables are measured at the end of the survey, and some respondents did not finish the full survey. Thus, the inclusion of these additional controls reduces the number of observations, which is the main reason why we exclude these variables from the baseline controls.

[60] To make the elicitation easier, in the prediction survey, we ask subjects to predict the effects of a 10 pp increase in the school share. In Figure 3, we divide those predictions by 10 to obtain the effect per 1 pp, so that it can be compared directly to the 2SLS estimates.

belief would have a negative effect on the protest rate for households *with* children (panel (a)) and a positive effect for households *without* children (panel (b)). They also are consistent with the mean of the experimental estimates in these two panels. However, the forecasts of most experts are inconsistent with the experimental results: most forecasts indicate effects that would be zero or that have the opposite sign compared to the experimental findings. Only a few expert predictions are close to the experimental estimates, even if we account for the sampling variation in the experimental estimates. More precisely, for households *with* children, only 41.1% of predictions are within the 90% confidence interval of the experimental estimate. For households *without* children, just 17.9% of predictions are within the 90% confidence interval of the experimental estimate.

At the end of the survey, we ask the experts to express how confident they feel about their forecasts. One notable finding is that experts do not feel confident about their predictions: on a scale of 1 to 5, where 1 is "not confident at all" and 5 is "extremely confident", the average confidence is 2.07.[61] In any case, we find that the comparison between the forecasts and experimental estimates is similar if we weight the forecasts by the confidence of the experts (results reported in Appendix B.9).

5.7 Non-Experimental Evidence

We complement this experimental evidence with some non-experimental evidence by including a survey question asking individuals to choose between hypothetical tax policies, in the spirit of Weinzierl (2014) and Saez and Stantcheva (2016). More specifically, we include a question about public school taxes. We present the respondent with a hypothetical situation in which two households (A and B) own homes worth $200,000 each. Both households are identical except that household A has two children enrolled in the public school district and household B has no children enrolled in the public school district. The respondent has to levy a total tax of $8,000, which can be spread across the two households in any way (e.g., assign all the burden to household A, all the burden to household B, or anything in the middle). According to the hypothesis of benefit-based taxation, respondents will want the household *with* children to pay more in taxes than the household *without* children, because the former benefits more from that government service. We find that most (58.8%) of the respondents behave according to the theory of benefit-based taxation, that is, they assign a higher tax burden to the household *with* children even though both homes are worth the same.[62] This evidence suggests that the logic of benefit-based taxation resonates with most taxpayers.

[61] More precisely, 25.0% of experts selected "not confident at all," 51.8% selected "slightly confident," 19.6% selected "somewhat confident," 3.57% selected "very confident", and 0% selected "extremely confident."

[62] For detailed results, see Appendix B.10.

A feature of property tax policy in the state of Texas is suggestive of benefit-based reasoning. Texas homeowners who are older than 65, most of whom do not have school-aged children, qualify for an exemption that limits their school taxes to the amount paid in the year that the owner turned 65, regardless of future increases in the home's proposed value (Texas Comptroller, 2021).[63] This exemption policy for households unlikely to have children is consistent with benefit-based reasoning.

6 Perceptions about Recapture

6.1 Accuracy of Prior Beliefs

Unlike the information on the school share, the information on recapture is not readily available in the Notice of Appraised Value from the DCAD. However, households may be informed about the recapture system through its media coverage. Homeowners also may deduce that recapture redistributes from more to less advantages districts; thus, knowing whether one lives in a more or less advantaged district may be enough to form a decent guess about their share of school taxes that are recaptured.

Figure 4(a) shows a histogram of the degree of misperceptions about the recapture share. The x-axis corresponds to the difference between the actual recapture share versus that perceived by respondents. A minority of subjects have accurate perceptions: around 20% of subjects guess the recapture share to be within $\pm$ 5 pp of the actual share. Misperceptions are significant in magnitude: the mean absolute error is 11.36 pp. However, the mean absolute error for the recapture share (11.36 pp) is somewhat less pronounced than that of the school share (16.57). The fact that misperceptions for the recapture share are somewhat smaller than those for the school share implies that there is less scope for the information provision experiment to update beliefs and thus less statistical power for the 2SLS estimates.

Unlike misperceptions about the school share, misperceptions about the recapture share have no systematic bias. Figure 4(a) shows roughly as many observations in the left half of the histogram (corresponding to overestimation) as in the right half (corresponding to underestimation). On average, subjects overestimate the recapture share by just 0.28 pp, as indicated by the mean error.

[63] The tax amount paid can increase if property improvements are made beyond maintenance and repairs. Homeowners also must apply to receive this benefit.

6.2 Belief Updating

Next, we summarize how subjects update their beliefs in reaction to the information provision about the recapture share. Figure 4(b) shows the results as a binned scatterplot. The x-axis corresponds to the gaps in prior beliefs, and the y-axis denotes the belief updating. As in the figure we used to show belief updating for the school share, the x-axis in Figure 4(b) shows the maximum revision we would expect if the respondents were to fully react to the provided information, and the y-axis shows the revision observed in practice. The red circles from Figure 4(b) correspond with subjects who are shown the feedback about the recapture share. Consistent with significant learning, there is a strong relationship between the belief revisions and prior gaps: an additional percentage point (pp) in perception gap is associated with a revision that is 0.632 pp higher. The gray squares from Figure 4(b) correspond with the subjects who do not receive information about the school share. In turn, the gray squares indicate a statistically significant (p-value<0.001) but economically small (0.099) degree of spurious revision. Most importantly, the degree of true learning corresponds to the difference in slopes between subjects who are shown the feedback and subjects who are not shown the feedback. This difference is large ($0.533 = 0.632 - 0.099$) and highly statistically significant (p-value<0.001). This difference suggests that a 1 pp information shock induces a 0.533 effect in posterior beliefs. Though large, this rate of information pass-through (0.533) is smaller than the corresponding rate for the school share (0.757).

The weaker effect of information shocks about recapture share, relative to school share, will result in less exogenous variation in posterior beliefs and thus less precisely estimated 2SLS coefficients. Many reasons help explain the weakly updated beliefs about recapture. For example, respondents may feel more confident in their prior beliefs about recapture, or they may trust the feedback on recapture less. Indeed, the recapture estimates are based on a number of assumptions, so subjects may naturally find the recapture feedback less persuasive. Last, subjects may pay less attention to the recapture feedback due to survey fatigue, as this information appears later in the survey.

6.3 2SLS Estimates

Let r_i^{post} be the posterior belief about the funds recaptured from individual i's own school district, in percentage points. Positive values indicate that individual i's district is a net contributor to the recapture system; in other words, $r_i^{post} = 40$ means that 40% of school taxes from household i's district are redistributed to disadvantaged school districts. Negative values, on the contrary, represent situations where individual i's school district benefits from recapture: $r_i^{post} = -30$ means that the school district receives an additional 30% over the amount of its own school taxes from taxpayers in other school districts.[64] We use the following econometric specification:

$$P_i^{2021} = \beta_0 + \beta_C^R \cdot C_i \cdot r_i^{post} + \beta_{NC}^R \cdot (1 - C_i) \cdot r_i^{post} + \beta_1 \cdot C_i + \epsilon_i \tag{22}$$

The two parameters of interest are β_C^R and β_{NC}^R for households *with* and *without* children, respectively. The benefit-based framework from Section 3 predicts that $\beta_C^R > 0$ and $\beta_{NC}^S = 0$. Again, the difference between these two parameters, $\beta_C^R - \beta_{NC}^R$, captures the benefit-based motivation behind public schools. As in the estimation of the change in perceptions about the school share, we estimate equation (22) using 2SLS to exploit the variation in r_i^{post} induced exogenously by the information provision experiment. As mentioned in Section 5.3, we estimate the effects of school share and recapture share jointly in the same 2SLS regression.

The 2SLS estimates for the recapture share are presented at the bottom panel of Table 2. In column (1) of Table 2, the dependent variable indicates if the subject protests directly in 2021. The coefficient for households *with* children is positive but close to zero (0.076) and statistically insignificant (p-value=0.875). For households *without* children, the coefficient is also positive (0.498) and large but borderline statistically insignificant (p-value=0.101). The difference between the coefficients for households with versus without children (0.076 and 0.498) is large but statistically insignificant (p-value=0.454).

Perhaps information about recapture share affects the intention to protest, but these effects do not materialize into behavior changes. To assess that possibility, column (2) of Table 2 is identical to column (1), except that it uses as a dependent variable the intention to protest instead of whether the household actually files a protest. As in column (1), the estimates from column (2) are all statistically insignificant. That is, the information on share recapture does not affect the intention to protest, even in the short term. The only coefficient from column (1) that is borderline significant (for households *without* children, with a p-value=0.101), is not even close to being statistically significant in column (2), and furthermore has the opposite sign.

One important caveat to keep in mind, however, is that the coefficients for the recapture share are less precisely estimated. For example, the standard errors of the coefficients for recapture share are more than twice the size of the corresponding standard errors for school share. As a result, the effects for recapture share should be more than twice as high as the effects of school share to have enough power to detect statistically significant effects. The less precise estimation for the coefficients for recapture share occurs for two reasons, both of which are difficult to anticipate ex-ante in the experimental design. First, as explained in

[64] The negative values can be lower than -100 because an ISD can receive more than 100% of the amount of its own school taxes in redistributed tax.

Section 6.1, the misperceptions about recapture share were smaller (mean absolute difference of 11.36 pp) than those about school share (mean absolute difference of 16.57 pp). Second, as documented in Section 6.2, conditional on a level of misperceptions, subjects updated their beliefs more strongly in response to the feedback about school share than in response to the feedback about recapture share.

To assess the robustness of the (null) effect of recapture share on protest behavior, we proceed with two falsification tests. Column (3) of Table 2 uses the dependent variable that indicates whether the household ever protested through an agent. As expected, the coefficients are statistically insignificant (p-values of 0.249 and 0.359 for households *with* and *without* children, respectively); the difference between the two coefficients (-0.207) is also statistically insignificant (p-value=0.486). For the second falsification test, column (4) of Table 2 uses the protest decision in a pre-treatment year (2020) as a dependent variable. As expected, the coefficients from column (4) (0.164 and -0.039 for households *with* and *without* children, respectively) are both statistically insignificant (p-values of 0.694 and 0.867); the difference between the two (0.203) is also statistically insignificant (p-value=0.664). We find similar results if we expand this falsification test to other pre-treatment years. Using binned scatterplots, we show that the (null) effects are not driven by non-linearities or outliers (results presented in Appendix B.8).

Table 3 presents additional robustness checks. Columns (1) and (2) of Table 3 reproduce the baseline specification given by columns (1) and (2) of Table 2 for reference, and columns (3) through (10) of Table 3 present the results under alternative specifications. In columns (3) and (4), we include additional control variables. In columns (5) through (8), we employ less stringent definitions of outliers. In columns (9) and (10), we exclude respondents who did not pass the attention check. The baseline results from columns (1) and (2) are consistent with the results from all the alternative specifications from columns (3) through (10).

We can also compare our 2SLS estimates to the expert predictions. However, as the experimental estimates are not estimated with sufficient precision, the results are not very informative (results presented in Appendix B.9).

In sum, we do not find any statistically significant evidence that perceptions about recapture affect the probability of filing a tax appeal. Although the null effects seem to conflict with predictions of the benefit-based framework, they must be considered in the context of limited statistical power. With those caveats in mind, one potential interpretation for the null effects would be respondents' altruism, as exemplified in the model extension presented in Appendix A. Intuitively, households may be willing to sacrifice government services if they believe those tax dollars will help the most disadvantaged communities.[65]

7 Conclusions

We present evidence that perceptions about where tax dollars go can have a significant effect on the willingness to pay those taxes. We conduct a high-stakes, natural field experiment with 2,110 homeowners who pay property taxes and have the opportunity to appeal their property tax assessment. We find that even though accurate information is publicly available and easily accessible, the homeowners have large misperceptions about how tax dollars are spent. Specifically, they do not know how much of their property taxes correspond with school taxes, and they misperceive the fraction of school taxes that are redistributed to school districts other than their own. The results from the field experiment are consistent with a basic prediction of benefit-based taxation. On the one hand, after learning that a higher share of property taxes funds public schools, households *with* children enrolled in public schools become less likely to appeal their property taxes. On the other hand, upon learning the same thing, households *without* children become more likely to appeal their property taxes. Beliefs about the school share have significant effects on the decision to file a protest, but beliefs about the recapture share do not have a significant effect, although this result must be considered in the context of its limited statistical power.

We implement our field experiment in Dallas Country, the second-largest county in Texas. Property taxes work almost identically in other counties in Texas and similarly across the country (Nathan et al., 2020), implying that our results are generalizable to other U.S. counties. Moreover, conducting the same field experiment in other U.S. counties would be straightforward. For instance, property taxes provide a significant source of school funding in most of the U.S. (Chen, 2021), and other states also redistribute property taxes across school districts similar to Texas' recapture system (Youngman, 2016). Finally, we provide detailed accounts of the implementation and data sources that other researchers can follow, and we are happy to share data, code, tips, and additional resources.

Though the main goal of this study is to provide causal evidence on the roots of tax compliance, our results also have some direct implications for tax and fiscal policies. First, our evidence suggests that governments may be able to design policies to increase their tax capacity. For example, governments may leverage transparency and accountability to improve perceptions about how tax dollars benefit taxpayers, which could boost tax compliance and increase taxpayer support of new taxes. We find evidence of large misperceptions about

[65] This interpretation is consistent with non-experimental evidence presented in Appendix B.10. Using a hypothetical question, we find that most households demand redistribution of school taxes across districts.

government spending, even when such information is publicly available. Governments thus may need to make an extra effort beyond just posting information on a website to reach constituents and educate them regarding where their tax dollars go. Relatedly, governments may want to simplify the connection between the taxes they collect and the government services provided, such as introducing new taxes that are earmarked for specific services. Local governments tend to do this well in that they typically break down property taxes into a school tax, a hospital tax, and so on. Even in the simple context of property taxes, however, we still find that taxpayers have large misperceptions about where their tax dollars go. In the case of state and federal governments, for which tax dollars follow a complicated path on their way to becoming public services, there is probably much room for improvement.

Our evidence suggests that normative considerations, such as benefit-based taxation, can have significant implications in practice. Modern tax theory largely disregards these types of normative considerations to focus on efficiency (Weinzierl, 2018).[66] However, taxpayers and politicians do not think like economists. The general public cares about efficiency, but they also care a lot about normative considerations, such as fairness. Taking these normative considerations into account can be helpful when designing efficient tax policies in a way that makes them more palatable to the general public. These normative considerations also may help clarify some puzzling observations indicating that some policies are chosen even though they are inefficient (Stantcheva, 2020).

References

Ajzenman, N. and R. Durante (2022). Salience and Accountability: School Infrastructure And Last-Minute Electoral Punishment. *Economic Journal, forthcoming.*

Alesina, A. and P. Giuliano (2011). Preferences for redistribution. In *Handbook of Social Economics*, Volume 1, pp. 93–131. Elsevier.

Athey, S. and G. W. Imbens (2017). The econometrics of randomized experiments. In *Handbook of economic field experiments, Vol. 1*, pp. 73–140.

Bergolo, M., M. Leites, R. Perez-Truglia, and M. Strehl (2020). What Makes a Tax Evader? *NBER Working Paper No. 28235*.

Bergolo, M. L., R. Ceni, G. Cruces, M. Giaccobasso, and R. Perez-Truglia (2021). Tax Audits as Scarecrows: Evidence from a Large-Scale Field Experiment. *American Economic Journal: Economic Policy, forthcoming*.

Blumenthal, M., C. Christian, and J. Slemrod (2001). Do Normative Appeals Affect Tax Com-

[66] Tax theorists have only recently explicitly integrated redistributive utilitarianism considerations into their models (Mankiw and Weinzierl, 2010; Scherf and Weinzierl, 2020).

pliance? Evidence From a Controlled Experiment in Minnesota. *National Tax Journal 54*(1), 125–138.

Bott, K. M., A. W. Cappelen, E. Ã. Sørensen, and B. Tungodden (2020). You've Got Mail: A Randomized Field Experiment on Tax Evasion. *Management Science 66*(7), 2801–2819.

Bottan, N. and R. Perez-Truglia (2020a). Betting on the House: Subjective Expectations and Market Choices. *NBER Working Paper No. 27412*.

Bottan, N. L. and R. Perez-Truglia (2020b). Choosing Your Pond: Location Choices and Relative Income. *The Review of Economics and Statistics, forthcoming*.

Cait, L., N. Jan-Emmanuel De, and Michael I. Norton (2018). The Power of Voice in Stimulating Morality: Eliciting Taxpayer Preferences Increases Tax Compliance. *Special Issue: Marketplace Morality*, 310–328.

Carrillo, P. E., E. Castro, and C. Scartascini (2021). Public good provision and property tax compliance: Evidence from a natural experiment. *Journal of Public Economics 198*, 104422.

Castro, L. and C. Scartascini (2015). Tax Compliance and Enforcement in the Pampas Evidence From a Field Experiment. *Journal of Economic Behavior & Organization 116*, 65–82.

Cavallo, A., G. Cruces, and R. Perez-Truglia (2017). Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. *American Economic Journal: Macroeconomics 9*(3), 1–35.

Chen, G. (2021). An Overview of the Funding of Public Schools. *https://www.publicschoolreview.com/blog/an-overview-of-the-funding-of-public-schools*.

Cruces, G., R. Perez-Truglia, and M. Tetaz (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics 98*, 100–112.

Cullen, J., N. Turner, and E. Washington (2020). Political alignment, attitudes toward government and tax evasion. *American Economic Journal: Economic Policy, forthcoming*.

Cullen, Z. and R. Perez-Truglia (2022). How Much Does Your Boss Make? The Effects of Salary Comparisons. *Journal of Political Economy 130*(3), 766–822.

Dallas Morning News (2018). Do your schools get your property tax dollars? *July 4, 2018*.

DellaVigna, S., D. Pope, and E. Vivalt (2019). Predict science to improve science. *Science 366*(6464), 428–429.

Dobay, N., F. Nicely, A. Sanderson, and P. Sanderson (2019). The best (and worst) of international property tax administration. Technical report, Council On State Taxation.

Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2021). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *Review of Economics and Statistics, forthcoming*.

Hoff, P. D. (2009). *A first course in Bayesian statistical methods*. Springer Science & Business

Media.

Huet-Vaughn, E. (2019, 2). Stimulating the Vote: ARRA Road Spending and Vote Share. *American Economic Journal: Economic Policy 11*(1), 292–316.

Huet-Vaughn, E., A. Robbett, and M. Spitzer (2019). A taste for taxes: Minimizing distortions using political preferences. *Journal of Public Economics 180*, 104055.

Imbens, G. W. and J. D. Angrist (1994). Identification and Estimation of Local Average Treatment Effects. *Econometrica 62*(2), 467–475.

Jones, P. (2019). Loss Aversion and Property Tax Avoidance. *Working Paper*.

Lockwood, B. B. and M. Weinzierl (2016). Positive and normative judgments implicit in U.S. tax policy, and the costs of unequal growth and recessions. *Journal of Monetary Economics 77*, 30–47.

Luttmer, E. F. P. and M. Singhal (2014). Tax Morale. *Journal of Economic Perspectives 28*(4), 149–168.

Mankiw, N. G. and M. Weinzierl (2010). The optimal taxation of height: A case study of utilitarian income redistribution. *American Economic Journal: Economic Policy 2*(1), 155–176.

McKenzie, D. (2012). Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics 99*(2), 210–221.

Musgrave, R. (1959). The Theory of Public Finance. *McGraw-Hill*.

Nathan, B., R. Perez-Truglia, and A. Zentner (2020). My Taxes are Too Darn High: Why Do Households Protest their Taxes? *NBER Working Paper No. 27816*.

Parker, W. and N. Friedman (2021). Zillow Quits Home-Flipping Business, Cities Inability to Forecast Prices. *The Wall Street Journal, November 2 2021*.

Saez, E. and S. Stantcheva (2016). Generalized social marginal welfare weights for optimal tax theory. *American Economic Review 106*(1), 24–45.

Scherf, R. and M. Weinzierl (2020). Understanding Different Approaches to Benefit-Based Taxation. *Fiscal Studies 41*(2), 385–410.

Seligman, E. R. A. (1908). Progressive Taxation in Theory and Practice. Technical Report 4.

Sinclair, M., J. O'Toole, M. Malawaraarachchi, and K. Leder (2012). Comparison of response rates and cost-effectiveness for a community-based survey: postal, internet and telephone modes with generic or personalised recruitment approaches. *BMC Medical Research Methodology 12*(1), 132.

Slemrod, J. (2019). Tax Compliance and Enforcement. *Journal of Economic Literature 57*(4), 904–954.

Stantcheva, S. (2020). Understanding Tax Policy: How Do People Reason? *NBER Working Paper No. 27699*.

Stock, J. H., J. H. Wright, and M. Yogo (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business and Economic Statistics 20*(4), 518–529.

Tax Policy Center (2021a). Amount of Revenue by Source. *https://www.taxpolicycenter.org/statistics/property-tax-revenue*.

Tax Policy Center (2021b). Amount of Revenue by Source. *https://www.taxpolicycenter.org/statistics/amount-revenue-source*.

Texas Comptroller (2021). Frequently Asked Questions. *https://comptroller.texas.gov/taxes/property-tax/exemptions/age65older-disabled-faq.php*.

Texas Education Agency (2021a). Excess Local Revenue. *https://tea.texas.gov/finance-and-grants/state-funding/excess-local-revenue*.

Texas Education Agency (2021b). Texas Public School Finance Overview: Biennium 2020-2021.

Texas Education Agency (2021c). What is House Bill 3? *https://tea.texas.gov/about-tea/government-relations-and-legal/government-relations/house-bill-3*.

U.S. Census Bureau (2021). Population, Dallas County, Texas. *https://www.census.gov/quickfacts/fact/table/dallascountytexas/POP010220*.

Villanueva, C. (2018). What is Recapture? *Center for Public Policy Priorities Report, August 30, 2018*.

Weinzierl, M. (2014). The promise of positive optimal taxation: normative diversity and a role for equal sacrifice. *Journal of Public Economics 118*, 128–142.

Weinzierl, M. (2018). Revisiting the Classical View of Benefit-based Taxation. *The Economic Journal 128*(612), F37–F64.

World Bank (2019). The Administrative Review Process for Tax Disputes: Tax Objections and Appeals in Latin America and the Caribbean.

Youngman, J. (2016). *A Good Tax*. New York: Columbia University Press.

Figure 1: Perceptions about the Share of Property Taxes Going to Public Schools

(a) Gap in Prior Beliefs **(b)** Belief Updating

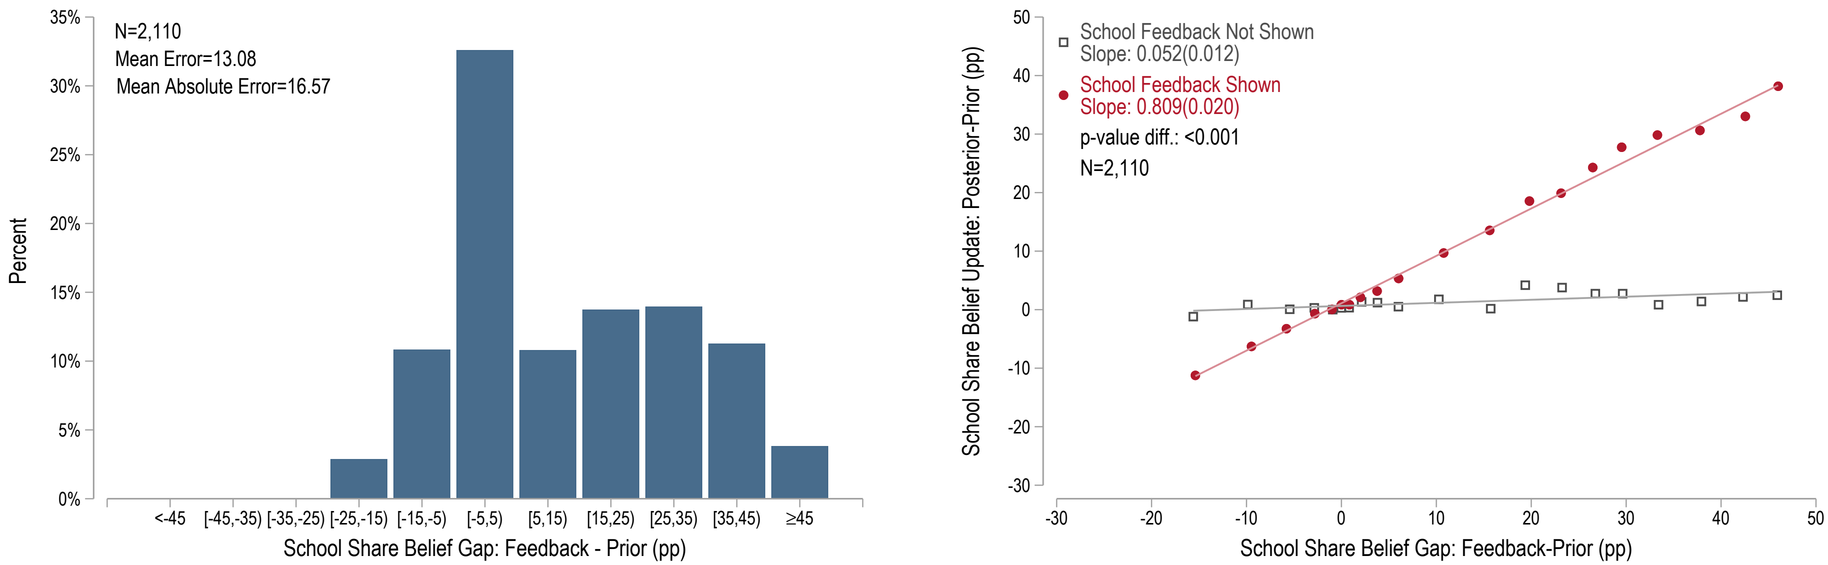

Notes: Panel (a) shows the gap in prior beliefs. The x-axis reports the difference between the actual school share and respondents' prior beliefs about the school share in 10 pp width bins. The y-axis reports the percentage of survey respondents in each bin. The upper left corner reports the total number of observations, the average error, and the average absolute error. Panel (b) shows how respondents update their beliefs using a binned scatterplot (using 20 equally sized bins). The x-axis reports the difference between the actual school share and respondents' prior beliefs about the school share. The y-axis reports the difference between posterior and prior beliefs (i.e., belief updating). Red circles (gray squares) represent the average update within each bin for the group of homeowners that were selected (were not selected) into the school share treatment. Each line corresponds to the fitted values from separate OLS regressions where the dependent variable is the actual update and the independent variable is the school share belief gap. The coefficients associated with the gap variable are reported in the upper left corner, as well as their robust standard errors (in parentheses), the p-value of the difference in the slopes, and the number of observations included in the analysis.

Figure 2: The Effects of School Share Perceptions on Protests: Additional Robustness Checks

(a) Event-Study Analysis **(b)** Binned Scatterplot (Reduced Form)

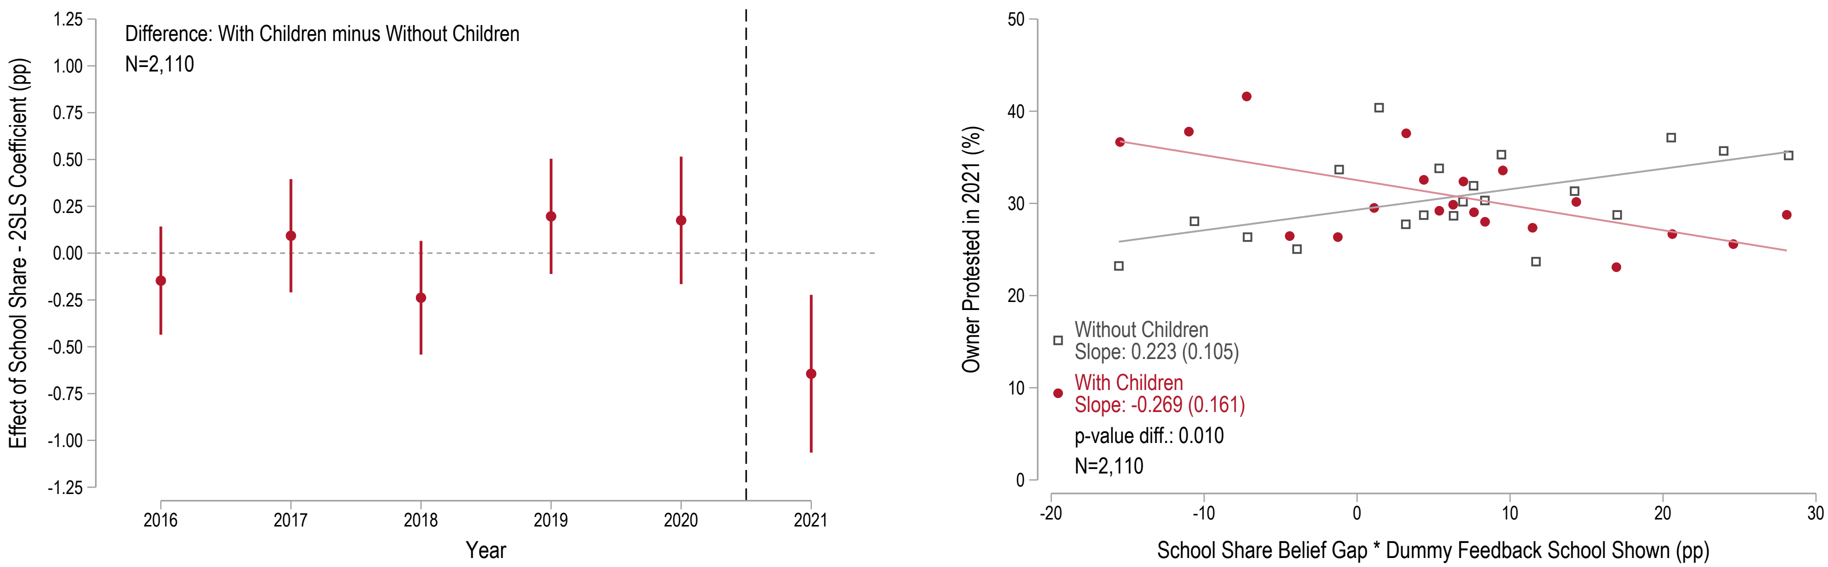

Notes: Panel (a) reports an event-study analysis of the differential effect of school share belief on the protest probability for households *with* children versus *without* children. The estimates plotted in this figure correspond with the 2SLS point estimate based on equation (21), with 90% confidence intervals based on robust standard errors. The coefficient plotted for 2021 is the coefficient reported in the "difference" row of panel (a), column (1) of Table 2. The remaining coefficients come from similar regressions but using the outcomes in pre-treatment years as falsification tests and restricting the pre-treatment controls to the corresponding years. The vertical dashed line separates the post-treatment year (2021) from the pre-treatment years (2016-2020). Panel (b) depicts a scatterplot representation of the reduced-form effect for households *with* and *without* children separately, using red circles and gray squares respectively and 20 equally sized bins. The x-axis corresponds to the interaction between the prior school share belief gap (defined as the difference between the actual school share and the prior belief about the school share) and a dummy variable that indicates if the homeowner was selected into the school share treatment group. The y-axis corresponds to the probability of a homeowner protesting directly in 2021. Each line corresponds to a separate OLS binned scatterplot regression, including the same control variables used in the 2SLS specification. The coefficients reported in the lower left corner and their (robust) standard errors are based on a unique regression that interacts the key variables with a dummy for having children at school (for the results in table form, see Table B.6). In addition we report the p-value of the difference in the effect for the two groups and the number of observations used in the estimation.

Figure 3: The Effects of School Share on Protests: Comparison to Expert Predictions

(a) With Children in Public School | **(b)** Without Children in Public School

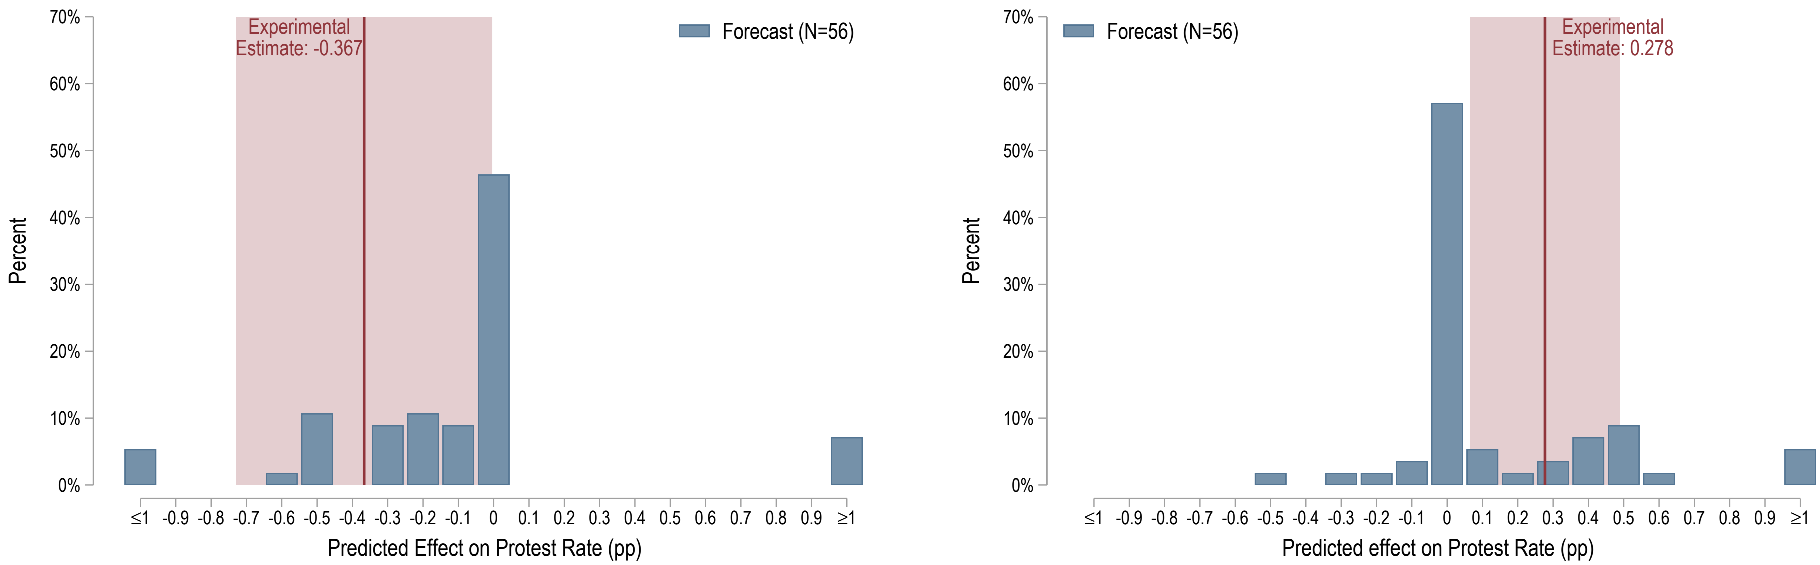

Notes: This figure shows the distribution of expert predictions about the effects of a 1 pp increase in school share beliefs on the probability that a homeowner fills a protest directly for households *with* children enrolled in the public school district (panel (a)) and households *without* children enrolled in the public school district (panel (b)), based on the data collected in the forecast survey. To make the elicitation easier, in the prediction survey we asked subjects to predict the effects of a 10 pp increase in beliefs about school share. For this figure, we divide those predictions by 10 and we obtained the effect per 1 pp to be able to compare these coefficients directly to the 2SLS estimates. In both panels, we pooled responses that were greater than 1 in absolute value into the extreme bins. The vertical red solid line corresponds to the experimental estimate based on the 2SLS specification reported in Table 2, panels (a) and (b). The pink shaded area corresponds to the 90% confidence interval. The full questionnaire for the prediction survey can be found in Appendix I.

Figure 4: Perceptions about the Share of School Taxes Affected by Recapture

(a) Gap in Prior Beliefs **(b)** Belief Updating

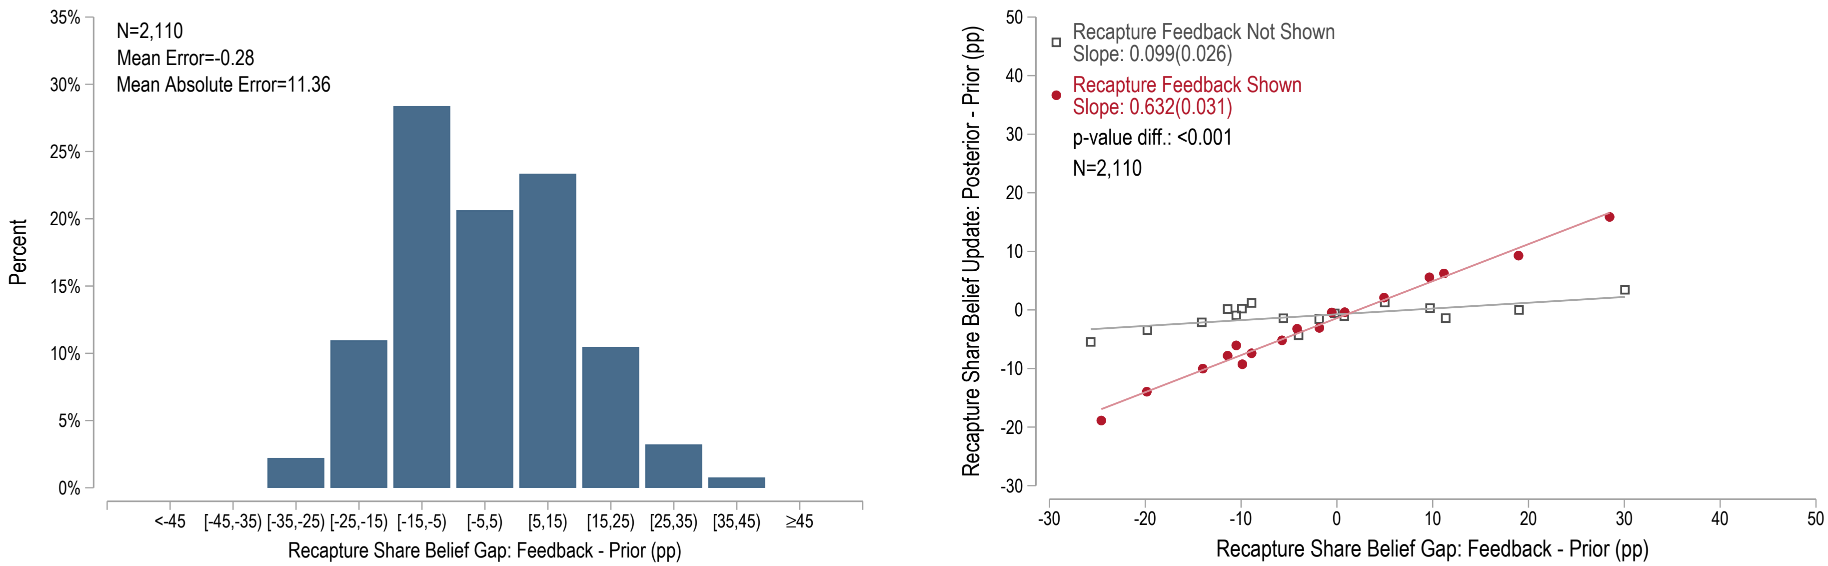

Notes: Panel (a) shows the gap in prior beliefs about the recapture share. The x-axis reports the difference between the actual recapture share and respondents' prior beliefs about the recapture share in 10 pp width bins. The y-axis reports the percentage of survey respondents in each bin. The upper left corner reports the total number of observations, the average error, and the average absolute error. Panel (b) shows how respondents update their beliefs using a binned scatterplot (with 20 equally sized bins). The x-axis reports the difference between the actual recapture share and respondents' prior beliefs about the recapture share. The y-axis reports the difference between posterior and prior beliefs (i.e., belief updating). Red circles (gray squares) represent the average update within each bin for the group of homeowners that were selected (were not selected) into the recapture share treatment. Each line corresponds to the fitted values from separate OLS regressions where the dependent variable is the actual update and the independent variable is the recapture share belief gap. The coefficients associated with the gap variable are reported in the upper left corner, as well as their robust standard errors (in parentheses), the p-value of the difference in the slopes, and the number of observations included in the analysis.

Table 1: Balance of Households' Characteristics across Treatment Groups

		Treatment Arm				
	All (1)	No Feedback (2)	Recapture Feedback (3)	School Feedback (4)	Both Feedback (5)	p-value test (6)
a. Admin. Records Variables:						
2021 Home Value ($1,000)	349.988	365.355	330.631	365.198	340.088	0.163
	(6.774)	(14.907)	(10.302)	(16.461)	(12.037)	
2021 Property Tax Amount ($1,000s)	7.738	8.018	7.448	7.960	7.546	0.292
	(0.129)	(0.296)	(0.218)	(0.287)	(0.228)	
School Share (%)	50.726	50.603	50.566	50.701	51.029	0.140
	(0.079)	(0.155)	(0.160)	(0.155)	(0.158)	
Recapture Share (%)	1.622	1.852	1.054	2.505	1.130	0.351
	(0.325)	(0.678)	(0.633)	(0.672)	(0.622)	
2020 Owner Protested (%)	18.057	23.121	14.815	19.883	14.684	0.000
	(0.838)	(1.852)	(1.530)	(1.764)	(1.527)	
2020 Agent Protested (%)	1.659	1.156	2.407	1.754	1.301	0.375
	(0.278)	(0.470)	(0.660)	(0.580)	(0.489)	
2019 Owner Protested (%)	13.365	15.029	10.926	14.035	13.569	0.238
	(0.741)	(1.570)	(1.344)	(1.535)	(1.478)	
2018 Owner Protested (%)	13.460	13.680	12.407	14.815	13.011	0.697
	(0.743)	(1.510)	(1.420)	(1.570)	(1.452)	
2017 Owner Protested (%)	10.853	11.561	11.111	11.891	8.922	0.400
	(0.677)	(1.405)	(1.354)	(1.430)	(1.230)	
2016 Owner Protested (%)	7.773	8.478	6.667	8.187	7.807	0.705
	(0.583)	(1.224)	(1.074)	(1.212)	(1.158)	
Multiple Owners (%)	24.645	22.929	24.444	25.146	26.022	0.693
	(0.938)	(1.847)	(1.851)	(1.917)	(1.893)	
Living Area (1,000s Sq. Feet)	2.313	2.317	2.302	2.331	2.302	0.959
	(0.022)	(0.046)	(0.042)	(0.046)	(0.040)	
Number of Bedrooms	3.428	3.432	3.398	3.423	3.459	0.609
	(0.016)	(0.032)	(0.033)	(0.034)	(0.031)	
Number of Baths	2.273	2.274	2.272	2.292	2.253	0.883
	(0.017)	(0.034)	(0.033)	(0.039)	(0.032)	
b. Survey Variables:						
With Children (%)	25.498	24.470	25.370	26.316	25.836	0.918
	(0.949)	(1.889)	(1.874)	(1.946)	(1.889)	
Female (%)	42.898	44.922	43.774	40.990	41.887	0.574
	(1.086)	(2.200)	(2.157)	(2.191)	(2.145)	
Age	49.608	49.711	49.381	50.438	48.945	0.146
	(0.234)	(0.470)	(0.481)	(0.461)	(0.460)	
Race: White (%)	44.300	44.727	47.818	44.422	40.265	0.103
	(1.092)	(2.200)	(2.178)	(2.220)	(2.134)	
Education: Grad. Degree (%)	38.309	39.844	37.761	38.446	37.240	0.841
	(1.069)	(2.166)	(2.114)	(2.173)	(2.104)	
Prior Belief: School Share (%)	37.642	37.741	37.186	37.935	37.726	0.918
	(0.394)	(0.804)	(0.760)	(0.790)	(0.800)	
Prior Belief: Recapture Share (%)	1.910	1.799	1.372	2.945	1.570	0.216
	(0.287)	(0.632)	(0.505)	(0.593)	(0.564)	
Observations	2,110	519	540	513	538	

Notes: This table lists pre-treatment characteristics averages. Statistics are based on the 2,110 homeowners that comprise the subject pool. Standard errors are reported in parentheses. The statistics in panel (a) are based on administrative records available at the DCAD website. The statistics in panel (b) are based on survey responses. Column (1) is based on the entire subject pool. Column (2) is based on homeowners not selected to receive any information. Column (3) is based on homeowners selected to receive information on the recapture share. Column (4) is based on homeowners selected to receive information on the school share. Column (5) is based on homeowners selected to receive information on both the recapture share and the school share. Column (6) reports the p-value of a test of equal means across the four treatment groups.

Table 2: Main Regression

	P_D^{2021} (1)	I^{2021} (2)	Falsification Tests P_A^{2021} (3)	 P_D^{2020} (4)
a. Effects of School Share:				
With Children	-0.367*	-0.408*	-0.028	0.110
	(0.221)	(0.234)	(0.118)	(0.181)
Without Children	0.277**	0.269*	-0.033	-0.065
	(0.129)	(0.144)	(0.051)	(0.097)
Difference (Children - No Children)	-0.644**	-0.678**	0.006	0.175
	(0.256)	(0.275)	(0.129)	(0.207)
b. Effects of Recapture Share:				
With Children	0.076	-0.313	-0.321	0.164
	(0.485)	(0.541)	(0.278)	(0.417)
Without Children	0.498	-0.101	-0.114	-0.039
	(0.303)	(0.325)	(0.124)	(0.234)
Difference (Children - No Children)	-0.422	-0.212	-0.207	0.203
	(0.563)	(0.620)	(0.297)	(0.468)
Cragg-Donald F-Statistic	30.10	30.22	30.10	30.02
Mean Outcome (Baseline):				
With Children	33.86	47.20	7.09	25.98
Without Children	28.83	44.87	4.08	22.19
Observations	2,110	2,090	2,110	2,110

Notes: Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. This table reports 2SLS estimates of equation (21) discussed in Section 5.3. Panel (a) reports the estimates corresponding to the school share treatment effect. We present the coefficients for households *with* children and households *without* children separately, as well as the difference between these two types of households. Panel (b) reports analogous results but for the recapture share treatment effects. The dependent variable in column (1) is an indicator variable that takes the value 100 if the subject protested directly in 2021. The dependent variable in column (2) is an indicator variable that takes the value 100 if the subject answered "very likely" to the question on the likelihood to protest in 2021. Columns (3) and (4) report the results of falsification tests. The dependent variable in column (3) is an indicator variable that takes the value 100 if the subject used an agent to protest in 2021. The dependent variable in column (4) is an indicator variable that takes the value 100 if the subject protested directly in 2020. Mean outcomes at baseline correspond with the mean of the dependent variables computed using the group of subjects who did not receive feedback about the school share nor recapture share.

Table 3: Main Regression: Robustness Checks

	P_D^{2021}	I^{2021}	P_D^{2021}	I^{2021}	P_D^{2021}	I^{2021}	P_D^{2021}	I^{2021}	P_D^{2021}	I^{2021}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
a. Effects of School Share:										
With Children	-0.367*	-0.408*	-0.429*	-0.457*	-0.330*	-0.250	-0.226	-0.088	-0.369	-0.418*
	(0.221)	(0.234)	(0.225)	(0.235)	(0.190)	(0.205)	(0.168)	(0.191)	(0.237)	(0.247)
Without Children	0.277**	0.269*	0.285**	0.286**	0.196*	0.321**	0.197*	0.256**	0.301**	0.324**
	(0.129)	(0.144)	(0.133)	(0.146)	(0.119)	(0.132)	(0.116)	(0.130)	(0.139)	(0.153)
Difference (Children - No Children)	-0.644**	-0.678**	-0.714***	-0.744***	-0.525**	-0.571**	-0.423**	-0.344	-0.671**	-0.743**
	(0.256)	(0.275)	(0.262)	(0.278)	(0.224)	(0.244)	(0.203)	(0.231)	(0.274)	(0.290)
b. Effects of Recapture Share:										
With Children	0.076	-0.313	0.141	-0.222	0.166	0.135	0.065	0.013	0.231	-0.059
	(0.485)	(0.541)	(0.478)	(0.536)	(0.417)	(0.451)	(0.330)	(0.373)	(0.442)	(0.492)
Without Children	0.498	-0.101	0.436	-0.125	0.414	-0.129	0.247	-0.051	0.473	-0.051
	(0.303)	(0.325)	(0.307)	(0.325)	(0.273)	(0.291)	(0.243)	(0.265)	(0.318)	(0.338)
Difference (Children - No Children)	-0.422	-0.212	-0.295	-0.098	-0.248	0.264	-0.182	0.063	-0.242	-0.009
	(0.563)	(0.620)	(0.559)	(0.616)	(0.500)	(0.536)	(0.394)	(0.438)	(0.527)	(0.579)
Cragg-Donald F-Statistic	30.10	30.22	29.68	29.68	35.26	35.34	47.35	47.55	34.35	34.35
Mean Outcome (Baseline):										
With Children	33.86	47.20	34.68	47.58	35.00	47.10	33.11	47.95	36.27	50.00
Without Children	28.83	44.87	29.12	45.10	29.77	45.64	29.53	46.09	29.33	44.28
Observations	2,110	2,090	2,070	2,070	2,335	2,309	2,482	2,454	1,807	1,807
Additional Controls			√	√						
2.5% Outliers					√	√				
1% Outliers							√	√		
Attention Check									√	√

Notes: * Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. This table reports 2SLS estimates of equation (21) discussed in Section 5.3. Panel (a) reports the estimates corresponding to the school share treatment effect. We present the coefficients for households *with* children and households *without* children separately, as well as the difference between these two types of households. Panel (b) reports analogous results but for the recapture share treatment effects. For reference, Columns (1) and (2) are identical to those in Table 2. The rest of the columns in this table use the dependent variables from Columns (1) and (2). Columns (3) and (4) add additional control variables collected in the survey: age, gender, college degree, and political party. Columns (5) and (6) drop 2.5% of the outliers at each tail of the distribution, instead of 5% which is the baseline specifications. Columns (7) and (8) drop 1% of the outliers at each tail. Columns (9) and (10) restrict the samples to subjects who passed the attention check included in the questionnaire. The attention check question can be found in the survey questionnaire in Appendix D. Each mean outcome corresponds with the mean of the dependent variable among subjects who did not receive feedback about the school share nor recapture share.