

NBER WORKING PAPER SERIES

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

Matias Giacobasso  
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 April 2023

We are thankful for excellent comments from Raj Chetty, Matthew Weinzierl, Austan Goolsbee, Steve Levitt, James Poterba, Dario Tortarolo, Sutirtha Bagchi and seminar participants at the NBER-Public Economics, University of Michigan, University of Chicago, University of Chicago-Advances in Field Experiments, RIDGE, IIPF, Journees LAGV, and NOVAFRICA. 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 could read about the hypotheses being tested), 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.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Matias Giacobasso, Brad C. Nathan, Ricardo Perez-Truglia, and Alejandro Zentner. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Where Do My Tax Dollars Go? Tax Morale Effects of Perceived Government Spending  
Matias Giacobasso, Brad C. Nathan, Ricardo Perez-Truglia, and Alejandro Zentner  
NBER Working Paper No. 29789  
February 2022, Revised April 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 measured how taxpayers perceive the destination of their tax dollars, such as the percentage of their property taxes that funds public schools. We find that even though accurate information is available, taxpayers still hold substantial misperceptions. We use an information-provision experiment to induce exogenous shocks to these perceptions. Using administrative data on property tax appeals, we measure the causal effect of perceived government spending on the willingness to pay taxes. We find that perceptions about government spending have a significant effect on the probability of filing a tax appeal and in a manner that is consistent with reciprocal motivation: individuals are more willing to pay taxes if they believe that the government services funded by those taxes will be of greater personal benefit to them. We discuss implications for the study of tax morale.

Matias Giacobasso  
Anderson School of Management - UCLA  
110 Westwood Plaza, C 3.10  
Los Angeles, CA 90077  
mggiacobasso@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

Why is tax compliance higher in some countries than in others? Why are some individuals more willing to pay their taxes than others? There are two schools of thought that offer potential explanations: institutions and tax morale. Abundant research shows that institutions have a large effect on tax compliance (Slemrod, 2019). For example, the introduction of withholding and third-party reporting caused a massive increase in tax compliance (Bagchi and Dušek, 2021). In contrast, little causal evidence shows that tax morale actually matters (Luttmer and Singhal, 2014). In this paper, we attempt to advance our understanding of tax morale by means of a natural field experiment, in a high-stakes context and via revealed-preferences.

Tax morale encompasses various potential mechanisms. We focus on one specific mechanism: our hypothesis is that individuals are more willing to pay taxes if they believe that the government services funded by those taxes will be of greater personal benefit to them. Our hypothesis is related to what Luttmer and Singhal (2014) calls *reciprocal motivation*: “the willingness to pay taxes in exchange for benefits that the state provides to them (...) even though their pecuniary payoff would be higher if they didn’t pay taxes.” Our hypothesis also relates to a normative principle known as *benefit-based taxation*, which can be briefly described as the “idea of basing tax liabilities on how much an individual benefits from the activities of the state” (Weinzierl, 2018). To test our hypothesis, we conducted an experiment to determine how taxpayers’ *perceptions* of how the government spends tax dollars affect their willingness to pay taxes.

Our experiment leverages the context of property taxes, which represents 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), nearly three times higher than the corporate income tax.<sup>1</sup> 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.

This context offers two key advantages to test our hypothesis of reciprocal motivation. First, our research design leverages the straightforward path between property taxes and the government services they fund, allowing to directly identify who benefits from what. For instance, households with children enrolled in local public schools benefit directly from publicly funded education, whereas households with no children enrolled in local public schools do not. For the sake of brevity, we refer to households with children enrolled in public schools as “households *with* children” and those without as “households *without* children.” The second

---

<sup>1</sup>For reference, the 2019 federal income tax generated \$1.717 trillion in revenue and corporate income tax generated \$230 billion (Tax Policy Center, 2021b).

advantage of this setting is that we can study the willingness to pay taxes via revealed preferences using households’ decisions to file property tax appeals, also known as tax protests (Nathan et al., 2020). Filing an appeal is a consequential, high-stakes action that households can take to reduce the amount they have to pay in property taxes.<sup>2</sup> In a nutshell, households can use the subjective nature of the appraisal process in their favor. If they feel like their taxes are too high, they can file a tax appeal to reduce their tax burden.<sup>3</sup>

We conducted a 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, to the extent that property tax appeals work similarly in other places, the results can be extrapolated to other settings. Dallas County is the second-largest county in Texas, with an estimated population of about 2.6 million in 2020 (U.S. Census Bureau, 2021) – indeed, Dallas County alone has a larger population than 15 of the 50 U.S. states. The county also is diverse along many dimensions, such as ethnicity, and has a relatively even distribution of Democrat and Republican supporters.<sup>4</sup>

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 between April and May of 2021, during which subjects could file a protest of their property taxes with the county. Our survey elicited whether the household has children enrolled in public schools to identify which subjects benefit directly from public school spending and which do not. We conducted an information-provision experiment a few weeks before households faced the opportunity to file a tax appeal. 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.

Our experimental design can be summarized as follows. First, we measure respondents’ perceptions about the share of their own property taxes that corresponds to school taxes and thus funds public school spending. For brevity, in the remainder of the paper we refer to this percentage as the household’s “school share.” The school share for the average house-

---

<sup>2</sup> When studying attitudes towards taxation, economists and other social scientists rely primarily on survey data. 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.

<sup>3</sup> For more details about how tax protests work, see the discussion in Section 2.3 and also Nathan et al. (2020) and Jones (2019).

<sup>4</sup> For example, in the 2012 presidential election, Barack Obama received 57% of the votes in Dallas County, whereas Mitt Romney received 42% (the remaining 1% of votes went to third-party candidates).

hold in Dallas County is about 49.78%.<sup>5</sup> 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. To illustrate, 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 reciprocal motivation and as noted in the randomized control trial (RCT) pre-registration, the expected effects of the information shock depend 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 principle of reciprocal motivation could have implications for tax redistribution. When taxpayers learn that their tax dollars are being spent in communities other than their own, they may be less willing to pay taxes because they do not receive benefits from the taxes they pay. We explore this additional hypothesis using a second treatment arm. Specifically, we leverage the significant redistribution of property taxes across school districts that occurs in some states. In Texas, for example, this redistribution is dictated by legislation often referred to by the media as the “Recapture Plan” or the “Robin Hood Plan.”<sup>6</sup> Thus, in the second treatment arm, we measure households’ perceptions about the share of their school funding that is redistributed away from, or towards, their own school district. 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 from an advantaged district is transferred to disadvantaged school districts.

We can measure the causal effects of the perceived recapture share using the information-provision experiment. According to the mechanism of reciprocal motivation, the belief about

---

<sup>5</sup> This average is calculated over 400,192 properties in Dallas County and excludes commercial properties and other non-owner-occupied residencies (for details, see Appendix B.1). Unless explicitly stated otherwise, all statistics about Dallas County are based on this sample.

<sup>6</sup> For the full history of property tax recapture in Texas, see for example Villanueva (2018).

recapture share should not affect the decision to file a tax appeal of households *without* children, because the diverted funds are being used for a service that does not benefit them directly anyways. By contrast, households *with* children should be more likely to protest upon learning that some of their tax payments are being diverted to other districts, because they were benefiting directly from the diverted funds.

The average subject in our sample owns a home worth \$349,988 and pays \$7,738 in annual property taxes.<sup>7</sup> 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%.<sup>8</sup> We also find significant variation in how the recapture system affects school districts in our sample, with some school districts diverting as much as 57% of their school districts' property taxes and others receiving as much as 23% additional funds from other districts.<sup>9</sup> 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. 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) protested directly in 2021. These tax protests are consequential. For instance, 65.4% of protests led 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 and find effects that are consistent with predictions of the framework of reciprocal motivation. 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 and economically significant. 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 sur-

---

<sup>7</sup> These estimated taxes are prior to any adjustments resulting from tax appeals.

<sup>8</sup> We show below that our school taxes information treatment is statistically significant for both of these types of households.

<sup>9</sup> The reported numbers refer to net transfers.

vey using a sample of 56 experts, most of whom are professors researching related topics. After receiving a brief explanation of the experiment, the experts are asked to forecast the experimental findings. Only a few of them were able to accurately predict the experimental findings. Most experts predicted that the beliefs on school share would have no effect on the likelihood of filing a tax appeal.

The results of the second treatment arm, about share of funds being recaptured, are unfortunately very imprecisely estimated and thus largely inconclusive. 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. As a result, the information shocks for the recapture share are not nearly as strong as those for the school share. Thus, the causal effects of the beliefs about the recapture share are very imprecisely estimated. It is important to note that the level of misperceptions and belief updating is difficult to anticipate before conducting the experiment. So, while ex-ante we expected to be well-powered for both treatment arms, ex-post we found out this was not the case for the second treatment arm. In an effort to mitigate publication bias (DellaVigna and Linos, 2022), we still report the analysis for the second treatment arm. Consistent with the hypothesis of reciprocal motivation, the belief about recapture share does not have significant effects on the decision to file a tax appeal among households *without* children – although this finding must be taken with a grain of salt due to the lack of sufficient statistical power. We do not find evidence of significant positive effects for households *with* children – however, the coefficient is so imprecisely estimated that we cannot rule out large positive effects.

Property taxes work almost identically in other counties in Texas and similarly across the country (Dobay et al., 2019; World Bank, 2019; Nathan et al., 2020).<sup>10</sup> These similarities imply that our results from Dallas County can be reasonably generalizable to other U.S. counties. Moreover, replicating our field experiment in other U.S. counties would be straightforward. Indeed, we propose the use of property tax protests as a novel context to study taxpayers’ preferences and tax compliance.<sup>11</sup> 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.

Our study relates and contributes to the literature on the role of tax morale in tax

---

<sup>10</sup> 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).

<sup>11</sup> One notable advantage of our setting is that it uses publicly available data, which facilitates replication efforts and avoids potential conflict of interests in partnerships with government organizations.

compliance decisions. Unlike the vast amount of causal evidence showing that institutions matter, there is little causal evidence showing that tax morale matters (Luttmer and Singhal, 2014). We contribute to this literature by providing novel evidence that tax morale can be a significant factor in practice. Moreover, we provide methodological innovations that other researchers can follow to better explore the role of tax morale.

Indeed, far from showing that tax morale matters for tax compliance, the existing causal evidence seems to suggest that its role is negligible. For example, a growing literature on correspondence experiments studies how moral suasion affects tax payments (Slemrod, 2019). These experiments consist of sending messages to taxpayers highlighting that paying taxes is the right thing to do and then measuring the effects of those messages on subsequent tax compliance. The moral suasion messages have been found to be largely ineffective (Slemrod, 2019). By comparison, messaging about institutions (e.g., audits, penalties) have been found to be highly effective. Based on this evidence, one natural interpretation is that tax morale is not important and that only institutions matter for tax compliance. Our results challenge this view: tax morale matters, but the existing correspondence experiments are ill-conceived to uncover the effects of tax morale.

Two innovations in our methodology allow us to shed light on tax morale, both of which are possible largely due to the novel research design linking data from a survey experiment to administrative tax data at the individual level. This approach is new to this stream of literature and rare even in broader economics research (Bergolo et al., 2020). First, moral suasion messaging in previous research typically has sought to affect individuals' tax morale by influencing individuals' preferences. However, such preferences are based on historical life experiences and may be too hard to change with a simple message (e.g., "it is important to contribute your part"). Instead of trying to influence preferences, we propose to study tax morale by inducing changes in beliefs. Indeed, a large literature shows that simple information-provision experiments can have significant and long-lasting effects on perceptions and expectations in a range of topics such as macroeconomic expectations (Cavallo et al., 2017) and salary perceptions (Cullen and Perez-Truglia, 2022).

The second innovation is our ability to measure heterogeneous effects via the linkage between the survey data and the administrative data. In the context of tax morale, there is scope for highly heterogeneous effects of information. As illustrated by our results on the school share, the same piece of information can have effects in opposite directions for different groups of subjects (i.e., households *with* children vs. households *without* children). It is possible that these large effects across different groups cancel each other out, on average, which would lead to the erroneous conclusion that tax morale is irrelevant for tax compliance. Using survey data to identify which households have kids enrolled in public schools and which



do not, we can measure the effects of the information separately for each group.<sup>12</sup> Another reason to expect heterogeneous effects relates to how subjects update their beliefs in response to new information. Households that underestimate their school share may adjust their beliefs upward when given accurate information, whereas households that overestimate their school share may adjust their beliefs downward when provided with the same information. Again, it is possible that these large effects across different groups cancel each other out, on average, which would lead to the erroneous conclusion that tax morale is irrelevant for tax compliance. Our survey allows us to measure prior and posterior beliefs, thus allowing us to fully elucidate the effects of the information.

Our findings are related to a few other studies. Some of the previously mentioned correspondence experiments are especially relevant because they include a treatment arm with a message related to the importance of taxes for the provision of community services (Blumenthal et al., 2001; Castro and Scartascini, 2015; Bott et al., 2020; De Neve et al., 2021; Bergolo et al., 2021).<sup>13</sup> As described above, our methodological innovations allow us to unpack effects that would have been otherwise hidden. Our study is also related to Cullen et al. (2020), who provide suggestive evidence of tax morale by showing that tax compliance can change with the partisan alignment of the government.<sup>14</sup> In another related study, 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 joint intervention decreased subsequent tax delinquency, although it is unclear whether the effects are due to the public recognition, the sidewalk, or both.<sup>15</sup> We contribute to this literature by disentangling a potential mechanisms at play: reciprocal motivation.<sup>16</sup> Beyond tax compliance, recent quasi-experimental evidence demonstrates how the salience of government spending

---

<sup>12</sup> Castro and Scartascini (2015) is a notable exception, as they provide direct evidence on treatment heterogeneity.

<sup>13</sup> The evidence is also mixed. 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. De Neve et al. (2021) find that information about government spending increases knowledge and appreciation about how taxes are spent but does not affect tax compliance. Castro and Scartascini (2015) find insignificant average effects.

<sup>14</sup> Huet-Vaughn et al. (2019) provide related laboratory evidence showing that the ideological match between the taxpayer and specific tax expenditures affects the willingness to pay taxes.

<sup>15</sup> Relatedly, Kresch et al. (2023) provide evidence showing that households with access to the city sewer system are more likely to pay property taxes.

<sup>16</sup> Also related to our study, Nathan et al. (2020) provide evidence that perceptions about the average tax rate affects households' decisions to file a protest. Although this result does not pertain to perceptions of government spending, it constitutes consistent evidence that fairness concerns play a significant role in the decision to file a tax appeal. Cait et al. (2018) provide evidence from a laboratory experiment showing that tax payments increase when participants have the opportunity to voice their preferences for how their tax dollars are to be spent.

can affect electoral outcomes (Huet-Vaughn, 2019; Ajzenman and Durante, 2022).

Our study also relates to a small but growing literature on the interplay between tax policy and normative considerations. The normative principle of benefit-based taxation was a prominent, and at times leading, approach among tax theorists in the early twentieth century (Seligman, 1908; Musgrave, 1959). However, it has been largely ignored by the modern optimal taxation literature, which instead focuses solely on efficiency aspects of taxation (Weinzierl, 2018; Scherf and Weinzierl, 2020). Moreover, a growing body of work seeks to incorporate other normative considerations into tax policy design (Mankiw and Weinzierl, 2010; Weinzierl, 2014; Saez and Stantcheva, 2016).<sup>17</sup> 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; Weinzierl, 2017). We contribute to this literature by providing the first revealed-preference evidence on this topic from a natural, high-stakes context.

We conclude by discussing some policy implications. Our evidence highlights the complexity of the use of transparency policies and communication strategies to boost tax compliance. For instance, messages sharing information on government services may have mixed effects on tax compliance based on whether the recipient of the message benefits from the advertised service. Indeed, our framework can explain the mixed results in the moral suasion literature, and our findings suggest some strategies that could raise average tax compliance.

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.<sup>18</sup> The second-highest component

---

<sup>17</sup> For instance, the normative considerations related to equality of opportunity or poverty alleviation.

<sup>18</sup> Variation in the school share across households ranges from 13.2% (1st percentile of the distribution) to 90.8% (99th percentile).

is the city tax (accounting for roughly 28% of property taxes), followed by hospital (10%), county (8%), college (4%), and special district (<1%) taxes.<sup>19</sup>

Dallas County has 14 main Independent School Districts (ISDs).<sup>20</sup> Homeowners who live within the geographical boundaries of a given ISD jurisdiction are subject to the tax rate for that ISD.<sup>21</sup> Households also have the right to send their children to the K–12 public school(s) 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.<sup>22</sup> 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.<sup>23</sup> 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.<sup>24</sup> According to data from the 2020 U.S. Census, about 90% of K–12 students in Dallas County attend a public school.<sup>25</sup>

## 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.<sup>26</sup> Due to the large tax amounts involved, the recapture system has been a topic of heated debate among politicians and the general public (Dallas Morning News, 2018). 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).<sup>27</sup>

In this paper, we focus on the *net* redistribution, which is the difference between the taxes recaptured by the state from the district (if any) and the amount distributed from that state

---

<sup>19</sup> See Appendix B.2 for more details.

<sup>20</sup> The total number of ISDs is sixteen, but two of them are extremely small and thus are excluded from the analysis. See Appendix B.1 for more details.

<sup>21</sup> School districts in Texas can set their own tax rates, but they must abide by certain state regulations. See Appendix B.2 for more details.

<sup>22</sup> For example, according to [www.GreatSchools.org](http://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).

<sup>23</sup> Charter schools are tuition-free public schools that receive funding directly from the state and do not receive funding from property taxes.

<sup>24</sup> Data accessed from <https://www.privateschoolreview.com/exas/dallas-county> on January 5, 2022.

<sup>25</sup> More precisely, 89% of kindergarten students and 92.5% of students in grades 1–12.

<sup>26</sup> 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.

<sup>27</sup> See Appendix B.2 for more details.

pool to the districts (for specifics on the recapture formula and this calculation, see Appendix B.2). 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.

## 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.<sup>28</sup> 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). After the notifications are sent, households have a month from the notification date to file a protest if they disagree with the proposed value. In 2021, the DCAD notified 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.<sup>29</sup> After reviewing the argument, the DCAD can (and often does) make an offer by mail or phone to reduce the assessed home value. If the homeowner refuses this settlement value or the DCAD does not offer a settlement, the appeal proceeds to a formal hearing with the Appraisal Review Board.<sup>30</sup> Once protests are resolved, the new tax amount becomes payable either immediately or at the billing date if it is later (i.e., 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.

---

<sup>28</sup> 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).

<sup>29</sup> 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).

<sup>30</sup> Homeowners can contest the Appraisal Review Board’s decision in court.

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 objectively “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 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 logic of reciprocal motives, 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 reciprocal motivation, we assume  $\gamma < 0$ : that is, when households benefit directly from government expenditures, they 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  $\alpha^S$  and  $\alpha^{NS}$  capture how households benefit from different types of expenditures. The parameter  $\alpha^S$  denotes how much a household *with* children enrolled in

public school benefits per dollar spent in public schools.  $\alpha^{NS}$  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:

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

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

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

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

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

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

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

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

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:

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

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,  $\alpha^S$ , 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 some of the main predictions 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.<sup>31</sup> 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:

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

$$B_{NC} = \alpha^{NS} \cdot NS \quad (12)$$

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

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

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

---

<sup>31</sup> The forces at play are similar from the opposite perspective, wherein a disadvantaged district receives funds from more advantaged districts.

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

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

$$\frac{\partial P_{NC}}{\partial r} = 0 \quad (16)$$

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

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

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

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

**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. We must keep in mind, however, that Prediction 2 no longer holds once we allow for altruistic taxpayers.

## 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.<sup>32</sup> As a second incentive, subjects were informed that survey respondents would be entered into a raffle for 20 prizes worth \$100 each.<sup>33</sup>

## 4.2 Survey Design

In this section, we summarize the main features of the survey.<sup>34</sup> 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 framework of reciprocal motivation, 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

---

<sup>32</sup> 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.

<sup>33</sup> 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.

<sup>34</sup> A sample of the full survey instrument is attached as Appendix H.

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 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-elicited 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, 2022).

The following module is about the recapture share.<sup>35</sup> 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

---

<sup>35</sup>Note that the recapture share is ISD-specific, whereas the school share is household-specific.

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).

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.)<sup>36</sup> When selecting this sample, we stratified the randomization at the ISD level to ensure wide representation of the beneficiaries and contributors to the recapture system.<sup>37</sup> 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

---

<sup>36</sup> For the full inclusion criteria, see Appendix B.1.

<sup>37</sup> For more details, see Appendix B.3.

accelerate delivery, we used a mailing company in Dallas County (i.e., the same county as all recipients). The mailing company dropped the letters off 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.<sup>38</sup>

Of 78,128 households invited to the survey, 2,966 answered the first two questions and 2,821 completed the two key modules on posterior beliefs about the recapture share.<sup>39</sup> 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 using 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).<sup>40</sup> 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, 2020), which 92.1% of respondents successfully passed. This passing rate is relatively high for a survey study, especially 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 to the DCAD’s records, had already hired a tax agent before starting our survey.<sup>41</sup>

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., 2022; Cullen and Perez-Truglia, 2022; Bottan and Perez-Truglia, 2020). For the baseline specification, we use a conservative definition of outliers that drops 467 subjects from

---

<sup>38</sup> More details about the timing of survey responses are provided in Appendix B.6.

<sup>39</sup> See Appendix B.3 for more details about the sample and Appendix B.6 for more details about attrition rates and balance tests.

<sup>40</sup> 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.

<sup>41</sup> For more details, see Appendix B.1.

the bottom 5% and top 5% of the distribution of prior misperceptions.<sup>42</sup> 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 results, such as event-study analyses.

Our subject pool self-selects to answer the survey for both households *with* and *without* children. The proportion of households *with* and *without* children who answered our survey, 25.5% and 74.5% respectively, approximately matches the proportion of families who have or do not have children in Dallas county, 32.3% and 67.4% respectively (Statistical Atlas, 2023).<sup>43</sup> 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.<sup>44</sup> 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.<sup>45</sup>

Columns (2) through (5) of Table 1 break down the average characteristics in each of the

---

<sup>42</sup> For more details on the distribution of outlier observations, see Appendix B.7.

<sup>43</sup> Moreover, our identification strategy will arise from comparing households *with* and *without* children and we discuss below how “salience” effects are likely to be similar for these two types of household.

<sup>44</sup> 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.1.

<sup>45</sup> 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).

four treatment groups. All characteristics shown in Table 1 are determined pre-treatment and thus should not be affected by the treatment assignment.<sup>46</sup> 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.<sup>47</sup> 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 are orthogonal to treatment assignment, which is expected given that subjects randomly receive treatment(s) 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 directly in 2021.<sup>48</sup> 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.<sup>49</sup>

Owners can file their own protests, which is the main focus of this paper. 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. Households also have the option to 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.<sup>50</sup> 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. First, because we provide information to the households and not to their agents, our experiment should not affect the agent’s behavior. Second, while it is possible that the information provided in our survey could influence whether a household hires an agent, that’s unlikely due to the nature of the

<sup>46</sup> 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.

<sup>47</sup> 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.

<sup>48</sup> The protest variable is based on data downloaded from the DCAD website on June 22, 2021.

<sup>49</sup> These calculations are based on data downloaded from the DCAD website in 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%).

<sup>50</sup> These statistics refer to 2021 protests in the control group.

setting. According to conversations with households, tax agents and representatives from assessor’s offices, households typically sign contracts with agents well in advance of the date when the proposed values are announced. Some households sign long-term contracts to file protests on behalf of the owner over many years. Last, consistent with the above arguments, in prior work we showed that a mail intervention had large effect on direct protests but no effect on protest through agents (Nathan et al., 2020). For all of these reasons, we study protests through tax agents separately, but in the spirit of a falsification test.

## 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 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 elicitation 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’ responses. Of these, 21.4% responded to the survey through the Social Science Prediction Platform, and the remaining 78.6% responded through our email invitation.<sup>51</sup> The final sample thus is comprised by 82.1% professors, 12.5% Ph.D. 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 Prior Beliefs

Transparency and accountability efforts have made information about property taxes publicly available. Each year, the Dallas Central Appraisal District (DCAD) provides homeowners in Dallas County 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

---

<sup>51</sup> 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.

taxes that funds public schools.<sup>52</sup> But 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, 2020).

Figure 1(a) shows a histogram of the degree of misperceptions about the school share.<sup>53</sup> 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 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  $\pm 5$  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).<sup>54</sup>

## 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., 2022).

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

<sup>52</sup> 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.

<sup>53</sup> 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.

<sup>54</sup> It might be thought that households *with* children have more accurate perceptions about the school share than households *without* children. We show in Figure B.4(a) that this conjecture is invalid since the distributions of perceptions are similar for these two groups.



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 shown feedback 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  $\alpha$  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.<sup>55</sup> This Bayesian updating model can be summarized by the following linear relationship:

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

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

---

<sup>55</sup> 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).

is consistent with other information-provision experiments that show evidence of spurious revisions (e.g., Fuster et al., 2022; Cullen and Perez-Truglia, 2022).

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

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

In this model, parameter  $\alpha$  represents true learning arising from the information provision (not spurious learning), whereas parameter  $\beta$  captures spurious learning. Parameter  $\alpha$  can be computed from the estimates in Figure 1(b). Specifically, the  $\alpha$  parameter corresponds to the difference in the regression slopes between the subjects who are and are not shown the feedback. Since  $\alpha$  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  $\alpha$  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 Econometric 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:

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

The term  $\epsilon_i$  represents the error. The two parameters of interest are  $\beta_C^S$  and  $\beta_{NC}^S$ . 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:

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

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  $(1 - C_i) \cdot T_i^S \cdot (s_i^{feed} - s_i^{prior})$ .<sup>56</sup>

We can illustrate 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. 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.<sup>57</sup>

The term  $X_i$  in equation (21) corresponds to a set of additional control variables. In

<sup>56</sup> 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 ( $T_i^S$ ).

<sup>57</sup> 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.

principle, the 2SLS model leverages the experimental variation, so control variables are not needed for causal 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.<sup>58</sup>

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 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.<sup>59</sup> 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 reciprocal motivation, 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 – for reference, this is roughly how much the average belief changed due to the information.<sup>60</sup> The estimates from column (1) of Table 2 indicate that this change

---

<sup>58</sup> 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).

<sup>59</sup> 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.

<sup>60</sup> More precisely, the average posterior belief increased by 10.69 pp due to the feedback (from 38.78 pp in the control group to 49.47 pp in the treatment group).

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), which is higher than the actual protest rate in the administrative data, 30.06%. For instance, a respondent may report a high protest likelihood in the survey, but then not protest due to filing frictions (Nathan et al., 2020).<sup>61</sup> 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,  $p\text{-value} < 0.001$ ).<sup>62</sup> Due to this imperfect correlation, the effects on the intention to protest at the time of answering the survey should not be 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 ( $p\text{-value} = 0.080$ ). 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 ( $p\text{-value} = 0.062$ ). Again, most importantly, the difference between the coefficients for households with children versus those without children (-0.408 and 0.269) is statistically significant ( $p\text{-value} = 0.014$ ).

---

<sup>61</sup> Spurious reactions (Cavallo et al., 2017) or other “salience” effects might explain part of the difference between the stated intention to protest in the survey and the actual protest. We do not think potential “salience” effects are an important concern in our setting. First, households protest days or weeks after answering the survey. Second, potential “salience” effects would presumably equally affect households *with* and *without* children and our identification strategy arises from comparing these two types of households.

<sup>62</sup> 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.

## 5.5 Robustness Checks

To probe the robustness of the school share results, columns (3) and (4) of Table 2 provide two falsification tests. For the first and most important falsification test, we exploit the timing of the information intervention in an event-study fashion. In column (3) of Table 2, we estimate the same baseline regression from column (1), except that we use as the dependent variable the protest decision in a pre-treatment year (2020), rather than in the post-treatment year (2021). Intuitively, since the information was provided in 2021, it could not possibly have an effect on the decision to protest as of a year earlier (2020). 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 households *with* and *without* children, respectively), precisely estimated, and statistically insignificant (p-values of 0.545 and 0.504); most importantly, the difference between households *with* children and *without* children is also close to zero (0.175) and statistically insignificant (p-value=0.398). Indeed, 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 in coefficients between households *with* children versus *without* children. Thus, the 2020 coefficient from Figure 2(a), which takes the value 0.175, is by construction equal to the corresponding coefficient from column (3) of Table 2. As expected, for each pre-treatment year (2016–2020) the coefficients are close to zero and statistically insignificant; by contrast, the coefficient is negative and statistically significant in the post-treatment year (2021).

The second falsification test uses a dependent variable that 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). 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).

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.<sup>63</sup>

The 2SLS model used for the results in equation (21) assumes a linear relationship between school share and the probability of protesting. This is 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. In a similar spirit, Appendix B.8 shows that the results are robust to an alternative approach that does not use 2SLS, which is simpler although it makes a less efficient use of the data.<sup>64</sup>

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 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). Note that 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. The results from columns (3) and (4) are similar to the baseline results from columns (1) and (2). If anything, the inclusion of the additional controls yields effects that are slightly stronger (-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

---

<sup>63</sup> For the detailed coefficients from the first-stage of the regression, see Appendix B.8.

<sup>64</sup> The appendix also shows the difference in sizes between the OLS regression not using the prior to posterior belief update and 2SLS regression using the belief update.

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 magnitude. 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.<sup>65</sup> 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 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: the majority of the forecasts predict either zero effect or an effect of 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

<sup>65</sup> 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.



are within the 90% confidence interval of the experimental estimate. That the majority of experts' predictions do not coincide with the experimental findings may not be surprising since their predictions are consistent with the general takeaway from the extant literature on how nudges affect tax compliance which suggests that deterrence nudges are effective whereas tax morale messages are less effective or have no effects whatsoever (see Bergolo et al. (2021)).

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.<sup>66</sup> 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 reciprocal motivation, 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 reciprocal mechanism, that is, they assign a higher tax burden to the household *with* children even though both homes are worth the same.<sup>67</sup> This evidence suggests that the mechanism of reciprocal motivation resonates with most taxpayers.

A feature of property tax policy in the state of Texas is suggestive of reciprocal motives. 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

---

<sup>66</sup> 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.”

<sup>67</sup> For detailed results, see Appendix B.10.

Comptroller, 2021).<sup>68</sup> 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. Also, it is probably widely known that the recapture system redistributes from more to less advantaged districts. As a result, if a homeowner knows whether he or she lives in a more or less advantaged district, that information on its own may be enough to form a decent guess about the recapture share.

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 substantially less pronounced than that of the school share (16.57). The fact that misperceptions for the recapture share are 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: on average, subjects overestimate the recapture share by just 0.28 pp. This can be seen directly from Figure 4(a), which shows that households are roughly equally likely to be in the left half of the histogram (corresponding to overestimation) as in the right half (corresponding to underestimation).

### 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. 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

---

<sup>68</sup> The tax amount paid can increase if property improvements are made beyond maintenance and repairs. Homeowners also must apply to receive this benefit.

practice. The red circles from Figure 4(b) correspond to 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 quite smaller than the corresponding rate for the school share (0.757).

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 that we use for the feedback 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. The most important implication of the weaker belief updating for recapture share (relative to school share) is that it will result in less exogenous variation in posterior beliefs and thus less precisely estimated 2SLS coefficients.

### 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.<sup>69</sup> We use the following econometric specification:

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

---

<sup>69</sup> 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.

The two parameters of interest are  $\beta_C^R$  and  $\beta_{NC}^R$  for households *with* and *without* children, respectively. The framework of reciprocal motivation 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 reciprocal 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 causal effects of the beliefs about the recapture share are very imprecisely estimated, so the results for this treatment arm are largely inconclusive. Consistent with the hypothesis of reciprocal motivation, the belief about recapture share does not have significant effects on the decision to file a tax appeal among households *without* children: the coefficient is positive (0.498) statistically insignificant (p-value=0.101). This finding must be taken with a grain of salt, however: since the coefficient is imprecisely estimated, we cannot rule out large effects, positive or negative.

To illustrate how imprecisely estimate this coefficient is, note that the standard error for recapture share is 135% larger than the corresponding standard error for school share (0.303 vs. 0.129). In other words, 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 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.

We do not find evidence of significant positive effects for households *with* children. The coefficient for households *with* children is positive (0.076) but statistically insignificant (p-value=0.875). Again, this coefficient is so imprecisely estimated that it does not really constitute evidence against the hypothesis of reciprocal motivation, because we cannot rule out very large positive effects. More precisely, the 95% confidence interval cannot rule out a positive coefficient of up to 1.02, which is several times the magnitude of the effects documented for the first treatment arm. Likewise, the difference between the coefficients for households with versus without children is statistically insignificant (p-value=0.454), but it is very imprecisely estimated so we cannot rule out large differences.

The coefficients from column (2) of Table 2 show that the results for recapture share are similar if we look at the intention to protest instead of the actual protest decision.<sup>70</sup>

The coefficients from column (3) and (4) of Table 2 show that, as expected, the results for recapture share do not show any effects on the falsification outcomes.<sup>71</sup> And Table 3 shows that the results for recapture share are similar under alternative specifications.<sup>72</sup> Some additional results are reported in the Appendix. Using binned scatterplots, we show that the results for the recapture share are not driven by non-linearities or outliers (results presented in Appendix B.8). We also compare our 2SLS estimates to the expert predictions.

## 7 Conclusions

Compared to abundant causal evidence on the importance of institutions for tax compliance, little causal evidence shows that tax morale is important. In this paper, we attempt to fill this gap by providing evidence from a natural field experiment. Our novel research design studies tax morale by linking data from a survey experiment to administrative tax records at the individual level. Our subjects are 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, households have large misperceptions about how tax dollars are spent. Through an information-provision experiment, we corrected misperceptions about where their tax dollars go. The effects of the information provision experiment are consistent with our hypothesis of reciprocal motivation. 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 and households *without*

---

<sup>70</sup> More precisely, column (2) of Table 2 is identical to column (1), except that the dependent variable is 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. The only coefficient from column (1) that is borderline significant, for households *without* children (p-value=0.101), is not even close to being statistically significant in column (2), and furthermore it has the opposite sign.

<sup>71</sup> More precisely, column (3) 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 (results presented in Appendix B.8). For the second falsification test, column (4) 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) and the difference between the two coefficients (-0.207) is also statistically insignificant (p-value=0.486).

<sup>72</sup> More precisely, 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: including additional control variables, using alternative definitions of outliers, etc. The baseline results from columns (1) and (2) are consistent with the results from all the alternative specifications from columns (3) through (10).

children become more likely to appeal their property taxes.

In this paper, we limited our exploration to the role of specific beliefs like the share of school taxes. However, we seek to make a more general methodological contribution: our research design can be used to study other mechanisms under the umbrella of tax morale. For instance, this approach could be used to assess the willingness to pay taxes in response to changes in the perceived quality of government spending or perceived corruption.

Our results stress the challenges of public communication policies. First, we document evidence of large misperceptions about government spending, even when such information is publicly available. For governments interested in educating their citizens on how tax dollars are spent, they should do more than post information on a website. Additionally, governments may want to simplify the connection between the taxes they collect and the government services those taxes support. Indeed, 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 how tax dollars are spent. 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 likely much room for improvement if the goal is to educate citizens.

Our experimental intervention was designed to disentangle causal mechanisms, not to increase average tax compliance. Nevertheless, our findings provide some hints for policymakers looking to boost tax compliance. Our results underscore the challenges and limitations of transparency policies and information campaigns. For example, a message highlighting a government service (e.g., public schools) can boost tax compliance among individuals who benefit most from that service (e.g., households *with* children), but it can reduce compliance from taxpayers who do not benefit from that service (e.g., households *without* children). As a result, these effects may cancel each other out, resulting in a null average effect on tax compliance. In some cases, this approach may even backfire. Our findings suggest that governments may be able to use reciprocal motives to boost average tax compliance, but only if they are willing to target the information (e.g., informing households *with* children about public school spending). Also, governments could try to persuade taxpayers that their tax dollars are spent efficiently or that their tax payments are not captured by corrupt politicians or wasted by bureaucrats. To the extent that these messages raise the average taxpayers' perception that their tax dollars are well spent, they also may increase the average tax compliance.

## References

- Ajzenman, N. and R. Durante (2022). Salience and Accountability: School Infrastructure And Last-Minute Electoral Punishment. *Economic Journal*, *forthcoming*.
- Athey, S. and G. W. Imbens (2017). The econometrics of randomized experiments. In *Handbook of economic field experiments*, Vol. 1, pp. 73–140.
- Bagchi, S. and L. Dušek (2021). The effects of introducing withholding and third-party reporting on tax collections: Evidence from the U.S. state personal income tax. *Journal of Public Economics* 204, 104537.
- 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. Giacobasso, 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 Compliance? Evidence From a Controlled Experiment in Minnesota. *National Tax Journal* 54(1), 125–138.
- Bott, K. M., A. W. Cappelen, E. Å. 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 (2020). Betting on the House: Subjective Expectations and Market Choices. *NBER Working Paper No. 27412*.
- Bottan, N. L. and R. Perez-Truglia (2022). Choosing Your Pond: Location Choices and Relative Income. *The Review of Economics and Statistics* 104(5), 1010–1027.
- Cait, L., J.-E. De Neve, and M. 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>.
- 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*.
- De Neve, J.-E., C. Imbert, T. Tsankova, and M. Luts (2021). How to Improve Tax Compliance? Evidence from Population-wide Experiments in Belgium. *Journal of Political Economy*, forthcoming.
- DellaVigna, S. and E. Linos (2022). Rcts to scale: Comprehensive evidence from two nudge units. *Econometrica* 90(1), 81–116.
- DellaVigna, S., D. Pope, and E. Vivaldi (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 (2022). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *Review of Economics and Statistics* 104(5), 1059–1078.
- 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*.
- Kresch, E. P., M. Walker, M. C. Best, F. Gerard, and J. Naritomi (2023). Sanitation and property tax compliance: Analyzing the social contract in Brazil. *Journal of Development Economics* 160, 102954.
- 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.



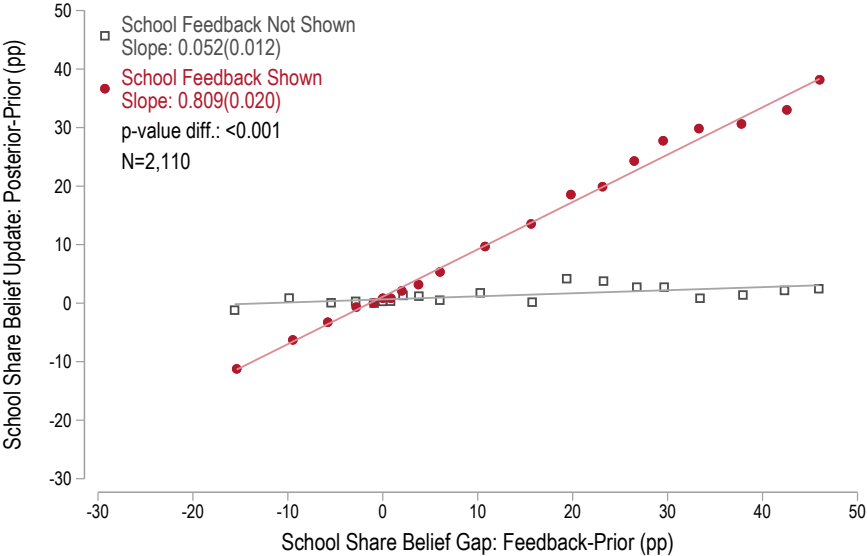

- Nathan, B., R. Perez-Truglia, and A. Zentner (2020). My Taxes are Too Darn High: Why Do Households Protest their Taxes? *NBER Working Paper No. 27816*.
- Parker, W. and N. Friedman (2021). Zillow Quits Home-Flipping Business, Cities Inability to Forecast Prices. *The Wall Street Journal*, November 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.
- Statistical Atlas (2023). The Demographic Statistical Atlas of the United States: Household Types in Dallas County, Texas. <https://statisticalatlas.com/county/Texas/Dallas-County/Household-Types>. Accessed: 2023-03-27.
- 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. (2017). Popular acceptance of inequality due to innate brute luck and support for classical benefit-based taxation. *Journal of Public Economics* 155, 54–63.
- 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

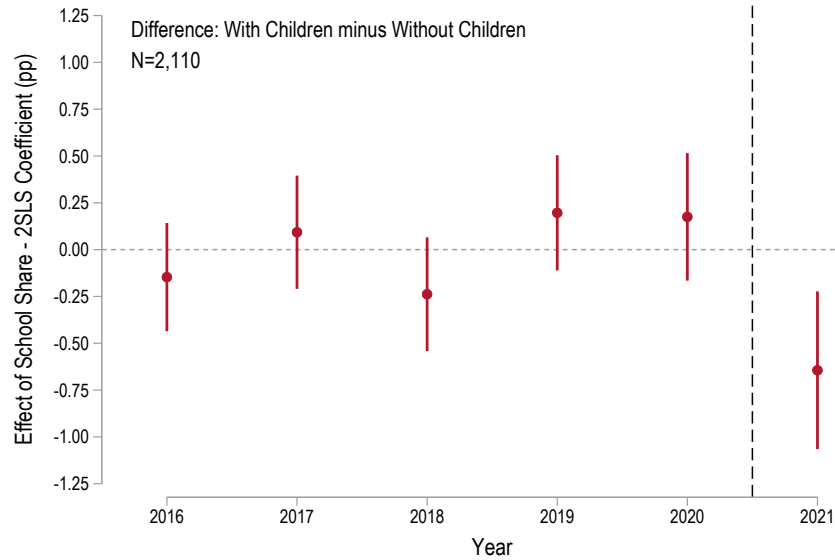


11


Notes: Panel (a) shows the gap in prior beliefs about the school share. 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 mean error, and the mean 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)

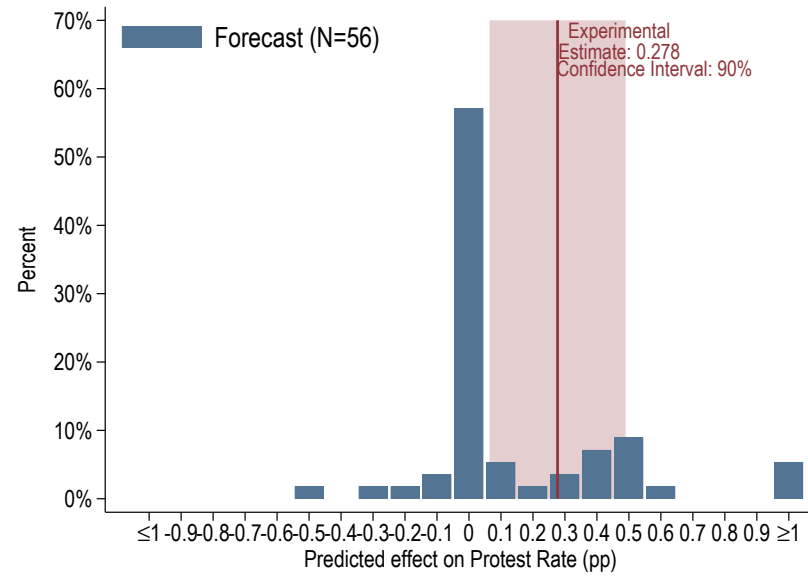
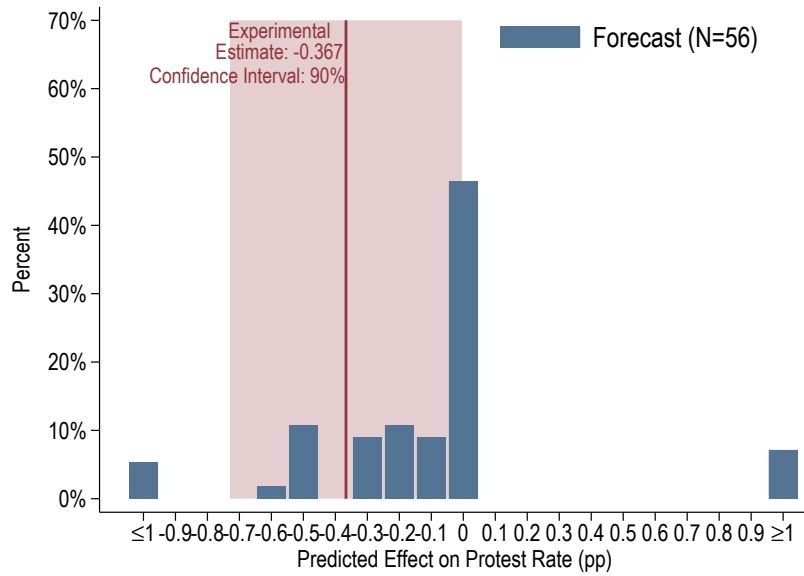


Notes: Panel (a) reports an event-study analysis of the differential effect of school share belief on the protest probability for households *with* children versus *without* children. The estimates plotted in this figure correspond with the 2SLS point estimate based on equation (21), with 90% confidence 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 direct protest 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

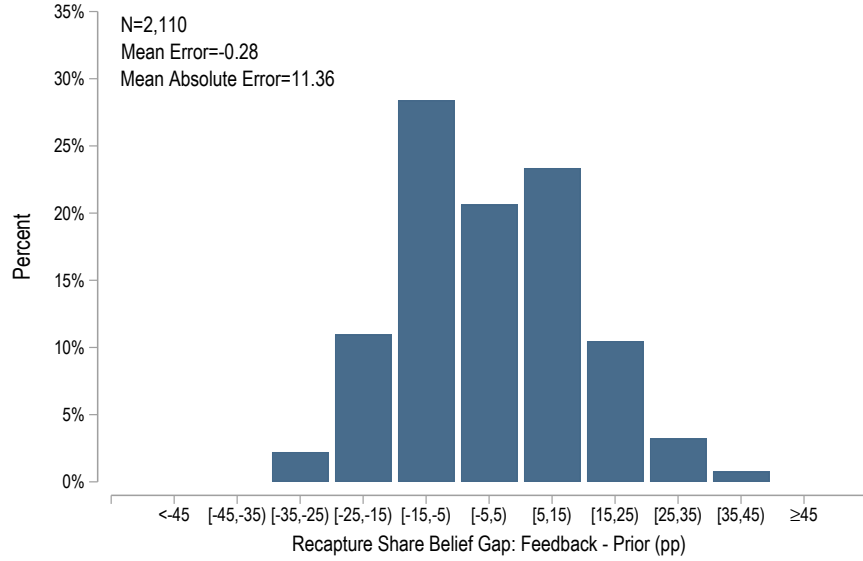
(b) Without Children in Public School



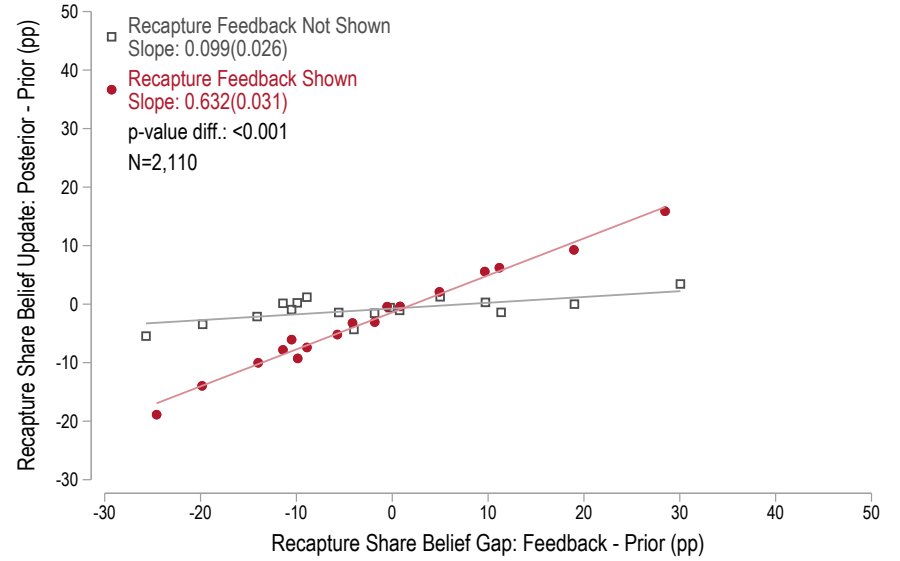
Notes: This figure shows the distribution of expert predictions about the effects of a 1 pp increase in school share beliefs on the probability that a homeowner files a protest directly for households *with* children enrolled in the public school district (panel (a)) and households *without* children 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 so these coefficients can be compared directly to the 2SLS estimates. In both panels, we pooled responses that were greater than 1 in absolute value into the corresponding extreme bins. The vertical red solid line corresponds to the experimental estimate based on the 2SLS specification reported in Table 2. The shaded area (in pink) 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 mean error, and the mean 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					p-value test (6)
	All (1)	No Feedback (2)	Recapture Feedback (3)	School Feedback (4)	Both Feedback (5)	
<b>a. Admin. Records Variables:</b>						
2021 Home Value (\$1,000)	349.988 (6.774)	365.355 (14.907)	330.631 (10.302)	365.198 (16.461)	340.088 (12.037)	0.163
2021 Property Tax Amount (\$1,000s)	7.738 (0.129)	8.018 (0.296)	7.448 (0.218)	7.960 (0.287)	7.546 (0.228)	0.292
School Share (%)	50.726 (0.079)	50.603 (0.155)	50.566 (0.160)	50.701 (0.155)	51.029 (0.158)	0.140
Recapture Share (%)	1.622 (0.325)	1.852 (0.678)	1.054 (0.633)	2.505 (0.672)	1.130 (0.622)	0.351
2020 Owner Protested (%)	18.057 (0.838)	23.121 (1.852)	14.815 (1.530)	19.883 (1.764)	14.684 (1.527)	0.000
2020 Agent Protested (%)	1.659 (0.278)	1.156 (0.470)	2.407 (0.660)	1.754 (0.580)	1.301 (0.489)	0.375
2019 Owner Protested (%)	13.365 (0.741)	15.029 (1.570)	10.926 (1.344)	14.035 (1.535)	13.569 (1.478)	0.238
2018 Owner Protested (%)	13.460 (0.743)	13.680 (1.510)	12.407 (1.420)	14.815 (1.570)	13.011 (1.452)	0.697
2017 Owner Protested (%)	10.853 (0.677)	11.561 (1.405)	11.111 (1.354)	11.891 (1.430)	8.922 (1.230)	0.400
2016 Owner Protested (%)	7.773 (0.583)	8.478 (1.224)	6.667 (1.074)	8.187 (1.212)	7.807 (1.158)	0.705
Multiple Owners (%)	24.645 (0.938)	22.929 (1.847)	24.444 (1.851)	25.146 (1.917)	26.022 (1.893)	0.693
Living Area (1,000s Sq. Feet)	2.313 (0.022)	2.317 (0.046)	2.302 (0.042)	2.331 (0.046)	2.302 (0.040)	0.959
Number of Bedrooms	3.428 (0.016)	3.432 (0.032)	3.398 (0.033)	3.423 (0.034)	3.459 (0.031)	0.609
Number of Baths	2.273 (0.017)	2.274 (0.034)	2.272 (0.033)	2.292 (0.039)	2.253 (0.032)	0.883
<b>b. Survey Variables:</b>						
With Children (%)	25.498 (0.949)	24.470 (1.889)	25.370 (1.874)	26.316 (1.946)	25.836 (1.889)	0.918
Female (%)	42.898 (1.086)	44.922 (2.200)	43.774 (2.157)	40.990 (2.191)	41.887 (2.145)	0.574
Age	49.608 (0.234)	49.711 (0.470)	49.381 (0.481)	50.438 (0.461)	48.945 (0.460)	0.146
Race: White (%)	44.300 (1.092)	44.727 (2.200)	47.818 (2.178)	44.422 (2.220)	40.265 (2.134)	0.103
Education: Grad. Degree (%)	38.309 (1.069)	39.844 (2.166)	37.761 (2.114)	38.446 (2.173)	37.240 (2.104)	0.841
Prior Belief: School Share (%)	37.642 (0.394)	37.741 (0.804)	37.186 (0.760)	37.935 (0.790)	37.726 (0.800)	0.918
Prior Belief: Recapture Share (%)	1.910 (0.287)	1.799 (0.632)	1.372 (0.505)	2.945 (0.593)	1.570 (0.564)	0.216
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 on the DCAD's 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 (control group). Column (3) is based on homeowners selected to receive information on the recapture share only. Column (4) is based on homeowners selected to receive information on the school share only. Column (5) is based on homeowners selected to receive information on both the school share and the recapture share. Column (6) reports the p-value of a test of equal means across the four treatment groups.

Table 2: Main Regression

	$P_D^{2021}$	$I^{2021}$	Falsification Tests	
			$P_D^{2020}$	$P_A^{2021}$
	(1)	(2)	(3)	(4)
<b>a. Effects of School Share:</b>				
With Children	-0.367*	-0.408*	0.110	-0.028
	(0.221)	(0.234)	(0.181)	(0.118)
Without Children	0.277**	0.269*	-0.065	-0.033
	(0.129)	(0.144)	(0.097)	(0.051)
Difference (Children - No Children)	-0.644**	-0.678**	0.175	0.006
	(0.256)	(0.275)	(0.207)	(0.129)
<b>b. Effects of Recapture Share:</b>				
With Children	0.076	-0.313	0.164	-0.321
	(0.485)	(0.541)	(0.417)	(0.278)
Without Children	0.498	-0.101	-0.039	-0.114
	(0.303)	(0.325)	(0.234)	(0.124)
Difference (Children - No Children)	-0.422	-0.212	0.203	-0.207
	(0.563)	(0.620)	(0.468)	(0.297)
Cragg-Donald F-Statistic	30.10	30.22	30.02	30.10
Mean Outcome (Baseline):				
With Children	33.86	47.20	25.98	7.09
Without Children	28.83	44.87	22.19	4.08
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, 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 subject’s protest likelihood 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 protested directly in 2020. The dependent variable in column (4) is an indicator variable that takes the value 100 if the subject used an agent to protest in 2021. 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 (the control group).



Table 3: Main Regression: Robustness Checks

	$P_D^{2021}$ (1)	$I^{2021}$ (2)	$P_D^{2021}$ (3)	$I^{2021}$ (4)	$P_D^{2021}$ (5)	$I^{2021}$ (6)	$P_D^{2021}$ (7)	$I^{2021}$ (8)	$P_D^{2021}$ (9)	$I^{2021}$ (10)
<b>a. Effects of School Share:</b>										
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>b. Effects of Recapture Share:</b>										
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. Columns (1) and (2) are identical to those in Table 2 (for reference). The rest of the columns in this table use the same 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 the 5% used in the baseline specification). 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 (see Appendix H for the survey). Each mean outcome corresponds with the mean of the dependent variable for subjects who did not receive feedback about the school share nor the recapture share (the control group).