NBER WORKING PAPER SERIES

MY TAXES ARE TOO DARN HIGH:
TAX PROTESTS AS REVEALED PREFERENCES FOR TAXATION

Brad C. Nathan
Ricardo Perez-Truglia
Alejandro Zentner

Working Paper 27816
http://www.nber.org/papers/w27816

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2020, Revised February 2021

We are thankful for excellent comments from John List, Michael Norton, Gil Sadka and other colleagues and seminar discussants. This project was reviewed and approved in advance by the Institutional Review Board at The University of Texas at Dallas. The experiments were pre-registered in the AEA RCT Registry (#0005992). Adrian Cadena Medina, Luisa Cefala, Dongwook Chun, Karl Dill and Santiago De Martini 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.

My Taxes are Too Darn High: Tax Protests as Revealed Preferences for Taxation
Brad C. Nathan, Ricardo Perez-Truglia, and Alejandro Zentner
NBER Working Paper No. 27816
September 2020, Revised February 2021
JEL No. C93,H2,H26,Z13

ABSTRACT

In all U.S. states, individuals can file a protest with the goal of legally reducing their property taxes. The decision to protest can provide unique revealed-preference evidence on individuals' support for taxation in a high-stakes, naturally-occurring context. We study the motives for protesting taxes using administrative records and two sources of causal identification: a quasi-experiment and a large-scale natural field experiment. We show that, consistent with selfish motives, the decision to protest is highly elastic to the private benefits and private costs of protesting. We find evidence of fairness motives too: consistent with conditional cooperation, a higher perceived average tax rate decreases both households' feelings of unfairness and their probabilities of protesting. Lastly, we study differences in protest choices between Democrats and Republicans.

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

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

A randomized controlled trials registry entry is available at https://www.socialscienceregistry.org/trials/5992
An online appendix is available at http://www.nber.org/data-appendix/w27816

Dedicated in loving memory of Alberto Alesina (1957–2020)

1 Introduction

Individual support for taxation plays a key role in modern democracies and is one of the oldest topics of study in political economy (Alesina and Glaeser, 2004). In this study, we introduce a novel context that allows us to study individuals' support for taxation via revealed-preferences in a high-stakes, naturally-occurring setting.

In all U.S. states, households can file protests with the goal of legally reducing the amount they have to pay in property taxes. These protests are consequential: while there is no guarantee that protests will reduce the tax bills, in practice they often do. The choice to protest taxes provides an interesting laboratory to study preferences for taxation via revealed-preferences. Filing a tax protest reflects a household's willingness to contribute to the provision of public services, holding constant what the tax rates are and how the tax revenues are being spent. We focus on this dimension of support for taxation, although this is of course not the only dimension that is important to a household's preferences for taxation. For instance, households may have preferences about what the tax rates should be and how tax revenues should be spent. We study this dimension of support for taxation via revealed-preferences and in a high-stakes, naturally-occurring context.

Although property tax protests are allowed across the United States (Dobay et al., 2019; World Bank, 2019), in this study we focus on one specific county (Dallas County, Texas) due to the logistical advantage of implementing the field experiment in a single location. However, there are lots of similarities in how tax protests work across the country.[1] Thus, the same methods that we apply in Dallas County could be replicated in other counties, and the main findings from Dallas County can be reasonably extrapolated to other U.S. counties as well. We chose Dallas County because it one of the largest and most diverse counties and where property taxes are especially important: since Texas does not have a state income tax, property taxes are a key source of revenue for the provision of government services. The average household from Dallas County is expected to pay around $5,916 in property taxes in 2020, corresponding to a tax rate of 2.01% on their homes' market value.[2]

The process to protest property taxes in Dallas County can be summarized as follows.

[1] The protest process is almost identical across all counties in Texas. There are, however, some differences across states. For example, owners must pay a filing fee to protest in Alaska and protests are less common in California because property taxes are based on taxable values that only change when the property is sold.

[2] These statistics are based on administrative data and focus on all single-family homes. Throughout this study, we use the term "tax rate" to refer to the household's effective tax rate (computed as the household's total property tax amount paid divided by its market value), rather than jurisdictional tax rates.

The Dallas Central Appraisal District (DCAD) formulates a proposed assessment of the property's market value and notifies the household of it (hereinafter, we refer to this amount as the *proposed value*). The property taxes will be calculated based on this proposed value. Next, the household can file a protest, arguing, for example, that the proposed value is too high.[3] Owners can either protest directly on their own (which is the main focus of this paper) or they can hire an agent to protest on their behalf. The protest is successful when the county lowers the effective assessed value. In 2020, a total of 8.40% of households in Dallas County filed a protest of their property taxes on their own, and an additional 8.42% of households protested with the help of an agent, resulting in a total protest rate of 16.83%.[4] We estimate that 69.7% of the direct protests in 2020 were successful and that these successful protests resulted, on average, in $485 in tax savings in the first year alone.[5]

In political economy, there are two classic channels that have been proposed to explain preferences for taxation and redistribution: selfishness and fairness. For instance, in the seminal model of Meltzer and Richard (1981), households are completely selfish and seek to maximize their own material well-being. On the other hand, more recent models propose that individuals do not only care about their own material well-being but they care about being treated fairly too (Alesina and Angeletos (2005); Benabou and Tirole (2006)). In this study, we combine quasi-experimental and experimental methods to study the role of selfishness and fairness considerations in the motivations to protest taxes. According to the selfish motives, taxpayers should be more likely to protest when they face higher private benefits or lower private costs. On the other hand, we study the role of fairness considerations in the form of conditional cooperation: i.e., whether taxpayers are less likely to protest if they think other taxpayers are contributing their fair share of taxes.

Teasing apart the motives for protesting taxes can be challenging. For example, even though there is no fee to file a protest and no risk of taxes going up when protesting, only a small minority of households choose to protest their taxes each year. One interpretation of this fact could be that the majority of households do not want to free-ride on the taxes paid by others. A different interpretation would be that most households want to free-ride, but they still choose not to protest because there are some hidden private costs from protesting, in the form of hassle costs, that are greater than their expected tax savings.

To study the role of selfish motives, we exploit exogenous variation in the private costs and benefits from protesting. For the private benefits, we exploit quasi-experimental variation

[3] We describe the protest process in more detail in Section 2 below.

[4] Approximately 34.5% of subject households protested at least once between 2015-2020, and 23.5% of subjects protested directly at least once during this time period.

[5] These estimates are based on data as of November 2020. As a result, this data does not reflect a few protests that might be resolved at a later date.

introduced by a facet of the Texas Property Code known as the homestead cap. All counties in Texas must use a cap when calculating taxes for households with homestead status.[6] This cap generates a sharp kink in the expected private benefits from protesting. When the proposed value is below the threshold, the tax amount would be reduced if the county were to revise the propose value downward. However, if the proposed value is above the threshold, a marginal reduction in the proposed value will have no effect on the tax amount. Exploiting this exogenous variation, we find that households are more likely to protest when they stand to gain more from protesting. A $100 increase in the marginal benefits from protesting causes an increase in the probability of protesting by 2.14 pp. This effect translates into an 11.26 elasticity – meaning that the decision to protest is highly elastic to the financial benefits.

To generate exogenous variation in the private costs from protesting, we designed and conducted a pre-registered field experiment. Any household can protest for free, so there are no pecuniary costs from protesting. However, we hypothesized that households could face significant hassle costs (Goolsbee, 2006; Benzarti, 2020; Sunstein, 2021; Benzarti, 2021). We conducted a mailing intervention aimed at reducing these hassle costs. We selected a subject pool of 78,462 households and sent letters to a random sample of 50,394 households with helpful information on how to file a protest. In the letters, we randomized the intensity of our aid provision. The basic letter treatment included information such as step-by-step guides for filling out the forms by mail or online. A second letter type included extra instructions based on the fact that, when protesting property taxes, one of the most challenging aspects of the process is preparing an argument to support the request. Further, the extra aid message included an argument tailored to each recipient, who could simply copy-paste it into their own protest form. These letters presented information about a comparison property near the recipient's own property that was similar in all observable characteristics and had been recently sold for a lower price than the market value proposed by DCAD.

We find that the letters had a large impact on the probability of filing a protest; and while the basic aid was helpful, the impact was significantly higher when the letter included the additional instructions. The letter with extra aid caused an increase in the protest rate of 4.98 pp (the intention-to-treat effect amounts to 3.51 pp); for comparison, the fraction of households protesting in the control group that did not receive a letter was 8.67 pp. In other words, a simple intervention by mail increased the protest rate by a whopping 57.4%. This evidence suggests that hassle costs are of first order importance in the decision to protest

[6] Homestead is a legal status that can be granted to properties that are the owner's (owners') primary residence. This legal status comes with several advantages related to property taxes (e.g., exemptions and caps) and other benefits such as exemption from forced sale for collection of debt.

taxes.[7] Moreover, we combine the experimental and quasi-experimental estimates to quantify the hassle costs. Our back-of-the-envelope calculations suggest that the hassle costs amount to \$232. Indeed, this estimate constitutes just a lower bound, as our intervention reduces some of the hassle costs but is far from removing them in full.

The finding that changes in private costs and benefits have a substantial effect on the decision to protest taxes suggests that selfish motives play a major role. One plausible interpretation is that most households would want to free-ride on their neighbors by protesting their taxes, but choose not to protest because the hassle costs are just too large relative to the expected savings. This point is perhaps best illustrated by a comparison between the average costs and benefits from protesting. On the one hand, for households who protested in 2020 the average tax savings was \$338 (including both successful and unsuccessful protests). On the other hand, our estimate of the costs from protesting, which is a lower-bound, sits at \$232. The fact that the average costs and benefits are in the same order of magnitude suggest that most households are probably weighting their private costs and benefits to decide whether to protest or not.

While the above evidence suggests that selfish motives are important for the decision to protest taxes, we provide evidence that fairness concerns play a role too. According to the conditional cooperation channel, taxpayers may be more willing to tolerate a higher tax rate if they think that the average household faces a high tax rate too.[8] This type of reciprocity has been documented in laboratory public good games (Gächter, 2007). However, it is not clear whether conditional cooperation would be significant in a context such as property taxes – among other things, the stakes are orders of magnitude higher for property taxes than for laboratory games, and individuals may have strong views about the government that do not manifest in laboratory games (Huet-Vaughn et al., 2019).

To test the conditional cooperation hypothesis, we included a second treatment arm in the field experiment. Our identification strategy leverages misperceptions about the average tax rate: households who pay below-average tax rates tend to underestimate the average tax rate, while household who pay above-average tax rates tend to overestimate the average tax rate. As a result, we can create exogenous variation in those perceptions by correcting the misperceptions through an information-provision experiment. In the letter sent to each household, we randomized whether we included information on the average tax rate paid in the county. We show that the information shocks induced by the experiment had a significant

[7] Our evidence of hassle costs is consistent with frictions found in other contexts such as the take-up of social benefits (Finkelstein and Notowidigdo, 2019) and filing income taxes (Bhargava and Manoli, 2015; Benzarti, 2020).

[8] Conditional cooperation is just one of the forms in which fairness concerns may arise. For instance, households may also feel more guilty about protesting if they are beneficiaries of the government services provided through the tax revenues, such as the public schools (Weinzierl, 2018).

effect on the probability of protesting, and in the direction predicted by the conditional cooperation channel: perceiving a higher average tax rate in the county causes a decline in the probability of protesting.

Moreover, we provide suggestive evidence that the effects of the information about the average tax rate is driven by fairness concerns. We conducted an online survey with a subsample of 1,888 subjects from the field experiment. We asked these taxpayers whether, relative to the county average, the tax rate that they face is fair or unfair. We show that the information provision experiment has a significant effect on perceived fairness and in the direction predicted by conditional cooperation: perceiving a higher average tax rate in the county causes taxpayers to see their own taxes as more fair. Furthermore, we provide evidence against alternative mechanisms, such as using the information about the average tax rate to make inferences about the protest process.

Our back of the envelope calculations suggest that conditional cooperation is not only statistically significant, but also significant in magnitude: the elasticity of the protest probability with respect to the perceived average tax rate is 1.63. The effects of conditional cooperation (elasticity of 1.63) are not nearly as strong as the effects of private benefits (elasticity of 11.26, reported above). On the one hand, this result may suggest that fairness concerns are not nearly as important as selfish considerations. However, conditional cooperation is just one of the possible manifestations of fairness concerns, so it is possible that fairness concerns add up to be as important as selfish concerns.[9]

The last part of the analysis assesses whether the motives for protesting differ along partisan lines. Survey data indicates that Republicans are much less likely to support income redistribution than Democrats (Di Tella et al., 2017; Alesina et al., 2020; Stantcheva, 2020). However, there is little evidence of whether those differences in survey responses materialize into differences in *behavior* – i.e., whether individuals put their money where their mouths are.[10] Our setting, thus, provides a unique opportunity to examine these partisan differences via revealed-preferences. We categorize households into Democrats and Republicans by matching the individuals from the property tax rolls to other sources of data such as the Texas voter files. We show that the probability of protesting as well as the motives behind the protests are qualitatively and quantitatively similar between Democrats and Republicans. Some of the differences along party lines are statistically significant, but they are always small in magnitude. This evidence suggests that some survey questions exaggerate the degree of partisan polarization in the support for taxation, and that there may be a lot more common

[9] Just to provide an example, some subjects may choose to protest because they consider it unfair that they pay a lot in taxes even though they do not send their kids to the local public school (Weinzierl, 2018).

[10] One notable exception is Cullen et al. (2020), who provide evidence suggesting that individuals evade less in federal taxes when the president in charge is from their same political party.

ground across party lines than previously thought.

This study contributes to various strands of literature. First, this study is related to literature on the decision to protest taxes. Protesting property taxes is allowed in all U.S. states and outside of the United States, too (Dobay et al., 2019; World Bank, 2019).[11] Despite being so common, tax protests have received little attention in the economics literature. To the best of our knowledge, there are two exceptions: Jones (2019) uses data on the decision to protest taxes to provide a test of loss aversion,[12] and Avenancio-León and Howard (2019) show that local governments place a disproportionate fiscal burden on racial and ethnic minorities and further document that some of these differences operate through tax appeals.[13] We contribute to this literature by identifying and measuring three important factors that affect the protest decision: private benefits, private costs, and conditional cooperation. Indeed, these findings can provide key inputs for the design of a more efficient and equitable system of tax appeals.

Second, we contribute to the literature on individual support for taxation and redistribution. Revealed preference evidence has proved elusive on this topic. For example, individuals do not get to choose how much they pay in taxes, thus making it impossible to learn about their preferences over tax rates from their choices. The typical way in which economists study preferences for taxation and redistribution is through survey data (Alesina and Giuliano, 2011; Cruces et al., 2013; Kuziemko et al., 2015), which has a number of well-known limitations – for example, individuals may say that they want high redistribution due to social pressure, but when the stakes are real they may act selfishly instead. A few efforts have been made to study support for taxation and redistribution from choice data, such as self-reported voting (Fisman et al., 2017; Epper et al., 2020) or charitable giving (Fisman et al., 2007) – however, these alternative approaches have limitations of their own.[14] We contribute to this literature by being the first to study support for taxation using the decision to protest taxes. While this context is not ideal, we believe it provides significant improvement in a number of

[11] In practice, protesting property taxes may be more common in some U.S. counties than in others for institutional differences (e.g., some counties charge significant fees to file a protest). Outside of the U.S., Dobay et al. (2019) found that protesting property taxes was allowed in all 10 countries they examined, and World Bank (2019) shows that property tax protests are allowed in several Latin American countries as well.

[12] In essence, Jones (2019) shows that the probability of protesting increases when the assessed value is revised upwards and decreases when the assessed value is revised downwards, but the effect is much larger (in absolute value) for the upwards revisions than for the downward revisions.

[13] Some studies look at property taxes more generally, without focusing on protests. One example is Cabral and Hoxby (2012), providing evidence of how the salience of property taxes can affect the equilibrium tax rates.

[14] For example, while income redistribution is certainly one issue that people care about when deciding whom to vote for, the decision can depend largely on non-economic factors such as the abortion and Second Amendment issues. Likewise, charitable giving can be largely driven by other factors such as religious participation (Bottan and Perez-Truglia, 2015).

respects. For instance, tax protesting is a type of high-stakes behavior that can be measured objectively with administrative data. Indeed, we provide detailed instructions so that other researchers can use this same experimental framework to study different research questions on this same topic.

Our study is also related to research on tax compliance (Luttmer and Singhal, 2014; Castro and Scartascini, 2015; Pomeranz, 2015; Hallsworth et al., 2017; Slemrod, 2019; Holz et al., 2020; De Neve et al., 2021). We contribute by studying a form of tax compliance that has been largely overlooked in the literature: tax appeals. Our study also relates to the literature on behavioral public finance showing that insights from behavioral economics can have important implications for public policy (Chetty, 2015). More precisely, we show that information frictions and conditional cooperation play a significant role in the decision to protest taxes.

The rest of the paper proceeds as follows. Section 2 describes the institutional context. Section 3 discusses the role of private benefits. Section 4 presents evidence on the private costs. Section 5 discusses the results on conditional cooperation. Section 6 analyzes the differences in tax protests by political party. The last section concludes.

2 Institutional Context

Dallas County is the second largest county in Texas with an estimated population of 2.6 million in 2020. In Texas, counties collect property taxes, which they use to fund various services, including schools, parks, roads, and the police and fire departments. The Dallas County tax assessor contractually collects property taxes. While the county collects property taxes on both residential and business properties, this study focuses on residential single-family homes. We use publicly available administrative data from the DCAD. For each home in the county, the data includes information on ownership, address, property characteristics (e.g., number of bedrooms), and historical yearly data on proposed and certified market values, exemption amounts, taxable values, tax rates as well as details on property tax protest records. Whenever needed, we complement the administrative records with other data sources.

The tax amount due is calculated by DCAD using a multi-step formula that starts with the county's proposed value of the home as of January 1st. Taxes due are a function of a host of factors such as the household's exemptions and the specific jurisdictional tax rates that pertain to the household, depending on the jurisdictions to which the home belongs.[15] Homeowners have the right to protest if they disagree with DCAD's proposed assessment value. Among other reasons, homeowners can protest if they believe the proposed value of their property is too high relative to the market value of comparable houses that were sold in the county, if their properties' proposed values are too high relative to the proposed values of comparable houses in the district, or if there are errors in the public records of the property (e.g., an incorrect number of bedrooms). For instance, according to the 2020 data for households that protested on their own, 91.87% of them selected the option "Value is over market value" in their online forms. When filing a protest, homeowners can also provide an "Opinion of Value", which is how much they believe their property was actually worth as of January 1st.

Homeowners can file protests on their own. We refer to this type of protest as *direct protests*. Because such protests are the focus of this paper, we always refer to this type of protests unless we explicitly note otherwise. Instead of protesting on their own, homeowners can hire an agent to protest on their behalf. In exchange for representation, agents normally charge some combination of a flat fee and a percentage of the tax savings (which can be as high as 50% of the tax savings). Protests through agents are less relevant to our study for a couple of reasons. Most importantly, we designed the field experiment specifically to reduce the hassle costs from protesting directly – which, if anything, should crowd out protests through agents. Second, the timing of the protests through agents makes it more difficult for them to be affected by the type of quasi-experimental and experimental variation used in our research design. According to anecdotal accounts, households often sign contracts with agents months before the proposed values are announced. Indeed, the decision to protest through an agent may have been made years ago, as agents offer long-term contracts to automatically protest on the owner's behalf every year.[16] For these reasons, the baseline specifications focus on direct protests and, for the sake of completeness, we report results for protests through agents in the Appendix.

The timing of the protest process is quite simple. Each year, the DCAD appraises the value of all homes in the county based on properties' market values as of January 1st. The DCAD shares the proposed values with homeowners through its website and, for most house-

[15] In Dallas County there are four county-level jurisdictions along with jurisdictions for each of the 23 of the cities, nine school districts, one Community College system, one hospital system, and 49 Public Improvement Districts (Source: `https://www.dallascounty.org/departments/tax/jurisdictions.php`). Appendix A.1.1 provides additional details about the property taxes in Dallas County and the data sources.

[16] While there is not any publicly available data on who entered into these long-term contracts, we do find some suggestive evidence in the protest data: households that protested through an agent in a given year have a high likelihood of protesting again through an agent in the following year. For instance, of the homeowners who protested through agents in 2019, 62.67% protested again through agents in 2020; in contrast, of the homeowners who protested directly in 2019, only 28.62% protested again directly in 2020.

holds, by mailing a “Notice of Appraised Value”. Households have a month from the notification date to file a protest. DCAD’s notifications include estimated taxes, which are based on each property’s proposed value. The term “estimated” is used to connote that, technically, property tax rates are determined later in the year, so the county uses the prior year’s jurisdictional tax rates to estimate taxes due in the Notice of Appraised Value. In practice, tax rate changes are uncommon so the approximation error is often negligible. In any case, these estimated taxes are the relevant object of study, as they represent the subjects’ expectations at their time of deciding whether to protest.[17] In 2020, DCAD presented the proposed values on May 15$^{\text{th}}$; as a result, the deadline to protest was June 15$^{\text{th}}$.

One key feature of this setting that is important for the interpretation of the results is that there is significant ambiguity in estimating market values. Because conducting full in-person appraisals is prohibitively expensive, the DCAD has to come up with its best guess for the market value of each property using statistical models and large datasets (e.g., recent home sales). The imperfections in these estimates are perhaps best illustrated by publicly-available data from websites such as Zillow.com and Redfin.com. When these companies publish estimates of the market value for the same property, their estimates tend to differ significantly, especially if that property has not been on the market recently. This ambiguity in market values leaves room for the owners to complain about the DCAD’s value assessments.[18] In a sense, households are not really “correcting” estimates that are obviously wrong. Instead, they are simply presenting a data point (e.g., the sales price of a neighboring home) to support their protest.

Based on our conversations with officers from some of the county appraisal districts in Texas, their prevailing view is that households use the subjective nature of the appraisal process not to complain about the county’s estimate of their home value *per se*, but simply as an excuse to complain about their taxes being too high. We provide suggestive evidence in support of this view: according to an independent estimate of household market value (Redfin), households still file protests when their properties have been under-assessed by the government.[19] This view is also consistent with household responses to open-ended questions

[17] For the sake of brevity, in this study we refer to *taxes*, but all analysis is technically based on *estimated taxes*. See Appendix A.1.1 for more details. The 2019 and 2020 jurisdictional tax rates were quite similar except for the city jurisdiction of Cockrell Hill (which dropped from 0.95% in 2019 to 0.85% in 2020) and Garland ISD (which decreased from 1.39% in 2019 to 1.26% in 2020).

[18] Households may have an advantage over the DCAD in that they know more about the specific attributes and condition of their own homes. On the other hand, households face a significant informational disadvantage in that they do not have access to the same models, data and expertise available to DCAD.

[19] Results reported in Appendix A.1.4.

from our survey mentioning that taxes are too high as their motivation to protest.[20]

Filing a protest directly is simple. Homeowners can protest using a paper form mailed by the DCAD to households that received a notification by mail because the proposed value increased relative to the previous year, a form from the Texas Comptroller that can be printed from the Internet, or via a simple online tool called uFile. To protest online, households need to search for their own name or address on a website, click on their account, and then follow some straightforward steps. In 2020, 75% of direct protests were filed online while the remaining 25% were filed by mail.

Protests can be resolved at different stages. Some protests are resolved because the owners accept the settlements proposed by the county. These settlements may be offered through informal channels, such as an email or phone exchange with a staff member from the DCAD. If an agreement is not reached, the protest advances to a formal hearing with a quasi-judicial entity called the Appraisal Review Board.[21] The formal hearing entails no risk: if the DCAD schedules a hearing and households do not attend, the protest is simply dismissed with no penalty.[22] After protests are resolved one way or another, the final assessed home values (from hereon, "certified" values) and tax amounts are calculated. Taxes become payable on October 1st, 2020 and, if unpaid, become delinquent after January 31st, 2021.

To aid in the interpretation of the results, we provide some basic descriptive statistics about the households in the sample. We focus on the sample of 423,607 single-family homes that were subject to property taxes in Dallas County in 2020, after excluding some potentially problematic cases such as households with missing data.[23] The average home in this sample has a value of $306,000 and pays $6,150 annually in property taxes. For a subsample of these subjects, we obtained individual-level demographic data from a private vendor.[24] The average subject is 52 years old, 65% are White, 9% are African-American and 20% are Hispanic. In different parts of the research design we focus on different subgroups of the main sample of 423,607 single-family homes, such as homes with homestead status or homes that were

[20] Our Field Survey (introduced in Section 4.2 below) included a final, open-ended question for subjects who wanted to share any thoughts with the researchers. Many households took this opportunity to mention their tax burden as a motivation for protesting.

[21] Formal hearings are typically conducted in person before a panel of three independent board members proposed by the DCAD and appointed by the Local Administrative District Judge of Dallas County. In response to the COVID-19 emergency in 2020, the DCAD staff did not conduct face-to-face negotiations, and all settlements were offered via email or telephone. Formal hearings were conducted over the phone with a single board member. If all else fails, the homeowner has the option to contest the decision in court.

[22] Of the 2020 direct protests that contain information on the form in which it was resolved, we find that 43.7% were settled informally, 35.2% were settled after a formal hearing, and 21.1% were either withdrawn or dismissed.

[23] For more details on the definition of this sample, see Appendix A.1.3.

[24] The company used the names and addresses to merge the records at the individual level. For more details about this data, see Appendix A.1.2.

selected to participate in the field experiment.[25]

In a given year, only a small share of households file a protest. In 2020, for example, 8.40% of homeowners in this sample protested directly.[26] If we include protests through agents too, the protest rate is 16.83%. This rate of protests has been quite stable in recent years: e.g., in the same sample of households from 2020, we find that 13.82% protested in 2017, 15.09% protested in 2018, and 13.89% protested in 2019. Even when looking at longer time horizons, it is still true that a minority of individuals file a protest: e.g., in the same sample of households from 2020, 23.5% protested directly (34.5% overall) at least once in the five-year period between 2015–2019.[27] The probability of protesting is not distributed evenly across households – for example, richer households and households who face higher tax burdens are significantly more likely to protest.

3 Private Benefits

3.1 Conceptual Framework

In this section, we use quasi-experimental variation in the pecuniary incentives to protest provided by Texas' property tax regulations. In Texas, homeowners may apply for homestead status for their primary residence. Among other benefits, the Texas Property Code guarantees that any increase in the appraised value of a homestead property is limited to 10% per year, which is referred to as the *homestead cap*. This regulation generates a sharp kink in the expected benefits from protesting. We exploit this kink as a natural experiment.

In practice, the amount of taxes that a household pays is calculated through a formula that involves the proposed value and the tax rates for the various jurisdictions within the county as well as other factors, such as the homestead cap and tax exemptions. Because households have the opportunity to protest every year, dynamic considerations may arise, too. For the sake of simplicity and to fix the intuition for the empirical analysis, however, we now introduce a simple model of the decision to protest. Let A be the proposed value of the household and T be the amount the household has to pay in property taxes. Under a simple proportional tax rate (τ), the tax burden without a homestead cap is the following:

$$T_{nocap} = \tau \cdot A \tag{1}$$

[25] In Appendix A.1.3 we provide descriptive statistics for each of the subsamples and used in the study, and show that they are roughly similar in a number of key respects.

[26] In all the analyses presented in this paper, we include protests that were marked as received by DCAD through July 15$^{\text{th}}$, 2020. For more details, see Appendix A.1.2.

[27] For more details, see Appendix A.1.4.

Let C denote the cost of protesting. Assume that households can protest ($P = 1$) or not ($P = 0$), and let $\Delta_A \geq 0$ be a random variable that corresponds to the reduction in A that would result from a protest. Then the expected net benefit from protesting is:

$$\mathbb{E}[U(P=1) - U(P=0)]_{nocap} = \tau \cdot \mathbb{P}(\Delta_A > 0) \cdot \mathbb{E}[\Delta_A | \Delta_A > 0] - C \tag{2}$$

and the household will protest if the above expected net benefit is positive and will not protest if it is non-positive. Now, let us introduce the homestead cap. Let the cap threshold be $\bar{A}$. Taking this threshold into consideration, the tax burden can be computed as follows:

$$T_{cap} = \tau \cdot \min\{A, \bar{A}\} \tag{3}$$

If the cap is not binding ($A < \bar{A}$), then T_{cap} is identical to T_{nocap}, and thus the decision to protest is not affected by the homestead cap. The interesting case is when the cap is binding ($A > \bar{A}$). As a result of a binding cap, the expected net benefit from protesting is as follows:

$$\mathbb{E}[U(P=1) - U(P=0)]_{cap} = \tau \cdot \mathbb{P}(\Delta_A > A - \bar{A}) \cdot \mathbb{E}[\Delta_A - (A - \bar{A}) | \Delta_A > A - \bar{A}] - C, \tag{4}$$

This equation can be re-arranged as follows:

$$\begin{aligned} \mathbb{E}[U(P=1) - U(P=0)]_{cap} =& \tau \cdot \mathbb{P}(\Delta_A > 0) \cdot \mathbb{E}[\Delta_A | \Delta_A > 0] - C \\ & - \tau \cdot \mathbb{P}(0 < \Delta_A < A - \bar{A}) \cdot \mathbb{E}[\Delta_A | 0 < \Delta_A < A - \bar{A}] \end{aligned} \tag{5}$$

Note that first two terms in the RHS in equation 2 are identical to the first two terms on the RHS in equation 5. Thus, the last term in equation 5 is the difference in incentives to protest introduced by the cap. The cap reduces the expected benefits from protesting when it is binding. Note that the expected benefits are lower the larger the difference between the proposed value and the homestead threshold ($A - \bar{A}$). The intuition is straightforward: absent a cap, a reduction in the assessed value will result in a reduction in the tax bill. When a household's proposed value is above the cap, however, the marginal reduction in the assessed value may not affect the final tax bill. If the proposed value is just $1 above the cap, then the first dollar reduction in the assessed value will not affect the tax bill but every dollar after that will. In that case, the cap should matter little to the household's decision to protest. However, if the proposed value is $15,000 above the cap, then none of the first $15,000 reduction in the assessed value will affect the tax burden, and thus the cap will substantially affect the expected benefits from protesting: the household will only see a reduction in the tax bill if the home value assessment is revised downwards by more than $15,000.

3.2 Main Results

Our analysis of the effects of the homestead cap on protest rates is based on the main sample of 423,607 single-family homes. About 74% of these households had a homestead status approved for 2020, and thus their homestead caps may be binding (and the remaining 26% of the households constitute the basis for a falsification test).

Figure 1 summarizes the main results. This figure is a binned scatterplot of the relationship between a given outcome and the distance between each household's proposed value (A) and its homestead cap threshold ($\bar{A}$). To be conservative, these baseline results are based on a narrow band around the homestead cap threshold ($15,000 above and below). In each panel, the horizontal axis measures the distance to the homestead cap threshold. The blue dots correspond to the households with proposed values below the homestead cap threshold, with the blue line corresponding to the linear fit. The coefficient from the linear regression is reported in blue too. For ease of exposition, we normalize all of the coefficients so that they correspond to the effects from a $10,000 increase in the proposed value. In turn, the red dots and red lines correspond to the households with proposed values above the homestead cap threshold. For the sake of simplicity, Figure 1 includes a minimal set of controls (the proposed value, a dummy for whether the household protested in the previous year, and a set of school district dummies).

The two panels on the left of Figure 1 (1.a and 1.c) correspond to the properties with homestead status – for which the homestead cap threshold can be binding. In turn, the two panels to the right of Figure 1 (1.b and 1.d) correspond to the properties without homestead status, for which the homestead cap threshold should be irrelevant and thus serve as a falsification test.

We start with Figure 1.a, in which the outcome variable measured on the vertical axis corresponds to the tax amount based on the 2020 proposed value.[28] The blue slope in the left half of the figure corresponds to properties right below the homestead cap threshold, while the red slope in the right half of the figure corresponds to the properties right above the threshold. Keep in mind that each of the slopes are just associations – it is only the difference between the two slopes that can be interpreted as a causal effect. As expected, there is a sharp kink at the threshold: after hitting the homestead cap threshold, households have to pay a lower tax amount than they would have needed to pay absent the homestead cap. This kink is not only large in magnitude, but also statistically significant: we can reject the null hypothesis that the slope to the left of the threshold (-97) is equal to the slope to the right of the threshold (-306), with a p-value$<$0.001.

[28] For an alternative measurement, Appendix A.2.1 reproduces the whole exercise but using the tax rate as a dependent variable instead of the tax amount.

The following thought experiment can illustrate how the homestead cap affects the marginal benefits from protesting. Consider a household with a proposed value that is \$10,000 above the homestead cap threshold – that is, imagine you are standing at +\$10K on the x-axis of Figure 1.a. The vertical gap between the red line and the dashed blue line, estimated at around \$209, corresponds to the tax amount that is capped.[29] What would happen if, due to a successful protest, the proposed value was reduced by \$10,000? In the presence of the homestead cap, the \$10,000 reduction in household value will not affect the tax amount of the household. That is, in presence of the homestead cap, the household will not benefit from the \$209 reduction in the tax bill, because the household was already benefiting from that reduction due to the homestead cap.

In sum, for a household sitting at +\$10K on the x-axis of Figure 1.a, the marginal benefit from protesting is lowered by \$209. If households care about the private benefits from protesting, we would expect the \$209 reduction in the marginal benefit from protesting would cause a reduction in the probability of protesting. To address that question, Figure 1.c is identical to Figure 1.a except that, instead of the tax amount, the vertical axis corresponds to the protest rate: i.e., an indicator variable that takes the value 100 if the owner protested directly in 2020 and 0 otherwise. We find that, as expected, there is a sharp kink in Figure 1.c at exactly the homestead cap threshold. This kink is not only large in magnitude, but also statistically significant: we can reject the null hypothesis that the coefficient to the left of the threshold (2.881) is equal to the coefficient to the right of the threshold (-1.604), with a p-value$<$0.001.

We can combine the results from Figure 1.a and Figure 1.c to quantify the magnitude of the effects of the homestead cap. Figure 1.a indicates that each additional \$10,000 above the homestead cap causes, on average, a reduction of \$209 in the tax amount. In turn, Figure 1.c indicates that being \$10,000 above the homestead cap causes a reduction in the protest probability of 4.485 percentage points. By taking the ratio of these two estimates, we conclude that for each \$100 reduction in the tax amount, the protest probability decreases by 2.14 pp. Moreover, we can also express this effect as an elasticity. The \$100 reduction in the tax amount corresponds to 1.9% of the average tax amount in the sample used for Figures 1.a and 1.c (\$5,153), while the 2.14 pp reduction in the protest probability corresponds to 21.4% of the average protest probability in this sample (9.98 pp). These percent changes suggest that there is a 11.26 ($= \frac{21.4}{1.9}$) elasticity of the protest probability with respect to the tax amount. In other words, the decision to protest is highly elastic to the private benefits from protesting.

[29] The red regression line predicts that when the proposed value is \$10,000 above the homestead cap the tax amount is \$306 lower. The dashed blue line predicts that when the proposed value is \$10,000 above the homestead cap the tax amount is \$97 lower. The difference between these two figures is \$209 (\$306 − \$97).

3.3 Robustness Checks

While the above analysis corresponds to properties that have a homestead status exemption (and thus are subject to the homestead cap), next we reproduce the analysis for properties that do not have a homestead status exemption (and thus are not subject to the homestead cap). This provides a sharp falsification test: we should not observe any kinks in the latter group – if there were kinks, that would suggest that the effects are not due to the homestead cap but due to some other confounding factor. The two panels in the right half of Figure 1 (i.e., 1.b and 1.d) correspond to the properties without homestead status, where the hypothetical homestead cap threshold is defined as 110% of the assessed value in the previous year (2019). As expected, we find that, in the right-side panels of Figure 1, there are no kinks at the homestead cap threshold: we cannot reject the null hypotheses that the coefficients are equal below and above the homestead cap threshold. Most importantly, the coefficients are precisely estimated in the right-side panels of Figure 1, meaning that we can rule out not only the large kinks shown in the left-side panel of Figure 1, but also rule out even small kinks.

For a second falsification test, we follow the logic of an event-study analysis. We reproduced the analysis from Figure 1.c, for the properties with homeastead status, but using as the dependent variable whether the household protested in 2019 (instead of whether the household protested in 2020). Whether the 2020 proposed value ends up being above or below the 2020 homestead cap should not affect whether a household protested a year prior, in 2019. If there was a significant difference, then this would suggest that the results from Figure 1.c are affected by a confounding factor. Moreover, we can extend this logic and reproduce the analysis not only for one year prior (2019) but for each of the other years we have data for (2015–2018). The results from this event-study analysis are presented in Figure 2.a. The rightmost coefficient (year 2020) corresponds to the effect of the homestead cap on the protest rate: i.e., the difference between the two slopes reported in Figure 1.c. The rest of the coefficients are estimated with an identical regression, except using a dependent variable that takes the value 100 if the household protested in 2015, 2016, 2017, 2018 and 2019, respectively.[30] As expected, the coefficients are consistently close to zero for each of the falsification years (2015-2019), and always highly statistically different from the coefficient for 2020.[31]

[30] Note that the set of control variables will not be identical: e.g., when the dependent variable is the 2020 protest choice, we control for whether the household protested in 2019, but when the dependent variable is the 2019 protest choice then we control for whether the household protested in 2018.

[31] Two of the falsification coefficients (for years 2015 and 2017) are borderline statistically significant, but still small in magnitude. Given that we estimate a total of 25 falsification coefficients in Figure 2, we should expect a few of them to be statistically significant just by chance.

In Appendix A.2.2 we present a number of additional robustness checks. We show that, in addition to affecting direct protests, the homestead cap affects the protests through agents and in the same direction, but those additional effects are smaller in magnitude as expected. In the baseline results we use a conservative specification based on a narrow band of $15,000 around the homestead cap threshold. We also show that the results are qualitatively and quantitatively similar if we use alternative bands. Last, we show that the homestead cap is consequential not only for the number of protests but also for the subsequent market values and tax amounts.

4 Private Costs

4.1 Conceptual Framework

The results presented above suggest that private benefits play a significant role in the decision to protest taxes. In this section, we explore the other side of the same coin: if households are responsive to private benefits, they should be responsive to private costs too.

Since there are no fees for filing a protest in Texas, the private costs that we are referring to are a form of non-pecuniary costs: the hassle costs. While filing a protest on your own is smooth in theory, it may be otherwise in practice. Some households may not even know where to start. Some may think the process is a lot more difficult than it actually is. Other households may be less sophisticated, and thus need step-by-step guidance on how to protest. Indeed, instructions on how to protest are not readily available. At the time of the experiment, only one official source online had instructions on how to file a protest: a PDF document posted on the Dallas CAD's website.[32] However, this document was long, had broad instructions, and was tucked deep into the DCAD's website. There were also a few unofficial online sources, such as blog posts, but those were usually incomplete, outdated, and difficult to find. Moreover, those sources often had a commercial interest, deliberately depicting the protest process as more complicated than it really is.

The hassle costs may be particularly large for one specific step in the protest procedure: providing an opinion of value for the home and an argument supporting it. While this process may be undertaken in different ways, according to anecdotal evidence, protesters typically identify a comparison property that has been sold recently for less than the proposed value of their own property. The comparable property's sale price can then serve as the opinion of value, and information about the recent transaction can be used as the argument. Finding

[32] This document can be found in the following address: `http://www.dallascad.org/Forms/Protest_Process.pdf`.

a proper comparison property entails a number of steps. The household needs to access a webpage, such as Zillow.com or Redfin.com, to identify properties that have been sold recently. The household needs to use the tools within the webpage to filter among recently sold properties with comparable features to theirs that were sold for less than the household's own proposed value within a few months of the start of the year. Although finding this comparison property manually may be relatively straight forward for some people, the task could be daunting for people with limited Internet access, limited Internet skills, or low financial literacy. Indeed, plenty of evidence indicates that households have trouble finding even easily accessible information such as information on the inflation rate or changes in average home prices (Cavallo et al., 2017; Bottan and Perez-Truglia, 2020a).

4.2 Experimental Design

We designed a mailing intervention aimed at reducing the hassle costs from protesting. If households care about their private costs from protesting, our intervention should increase the probability that they file a protest.

Subjects can be randomly assigned to receive no letter or a letter. Figures 3 and 4 show the first and second page of this sample letter, but with the addition of some red boxes highlighting the parts that were randomized.[33] The letters were sent on behalf of researchers at The University of Texas at Dallas and included a number of measures aimed at showing that they came from a legitimate source.[34] The letters were tailored to the recipients too: the salutation at the top of the first page included the name of each recipient; and their names and addresses were shown at the bottom of the second page (which appeared through the envelope window).[35]

In the first treatment arm, subjects were randomized to receive one of two types of letters: the *basic aid letter* or the *extra aid letter*. The *basic aid letter* consisted of a number of useful tips to help the recipient file a protest. All of that information is found on the *first* page of the letter, a sample of which appears in Figure 3. A key part of the first page is that it mentions recipients could find instructions on how to file a protest on the project's website.

[33] Appendix C provides a full-page sample of the letter without the additional red boxes.

[34] The envelope (see Appendix B for a sample) includes the logo of The University of Texas at Dallas, a well-known institution in Dallas County, and the name of one of the professors from that university. The envelope also included non-profit organization postage. The letter contained the official logo of The University of Texas at Dallas in the header as well as a physical address that they could write to and the URL for the study's website. The website included general information about the study (without discussing any hypotheses or what the study was about) as well as contacts for the researchers and the Institutional Review Board.

[35] Some properties are jointly owned by multiple individuals (typically, husband and wife). In those cases, we sent a single letter addressed to all the individuals listed as owners.

We designed our website instructions to be concise, easy to follow, and as explicit as possible. Appendix D shows screenshots of the entire website, which included step-by-step instructions on how to file a protest online or by mail. These walkthroughs included hyperlinks to the relevant websites as well as screenshots of a protest using the information from a fictitious household for added clarity.

The *extra aid letter* is identical to the basic aid letter except that it includes additional information on the second page that is aimed at providing even further guidance on how to protest. Figure 4.a shows what the second page of the letter would look like if the recipient was assigned to the basic aid treatment, while Figure 4.b shows what it would look like under the extra aid treatment. The extra aid message in Figure 4.b is highlighted inside of a red box with dashed lines (this box is shown for expositional purposes and was not included in the actual letters sent to subjects). The extra aid message also provides information to reduce the hassle costs created by the need to provide an opinion of value and an argument for supporting it, as described above.

The first paragraph of the extra aid message began by providing some facts about the protest filing process, such as that protesting is simple, can be done without an agent, may not require a hearing (which could be intimidating to some subjects) when the DCAD proposes a settlement offer; and even if a hearing is scheduled, there is not any risk if it is not attended. Then, it provided an argument for the protest. More specifically, we presented the most common type of argument: based on the recent sales price of a comparable property, the proposed value for the property is over the market value.[36] To make it simpler to use this information, we presented it much how it would look on the actual protest form: with a check mark in the "Value is over market value" box, an opinion of value field with the sales price of the comparison property filled in, and then a handwritten note with a usable argument. For example, in the sample letter shown in Figure 4.b, the handwritten note reads, "I found a home that is similar to mine but was recently sold for less than my home's appraised market value. The property located at 2234 Meadowstone Dr. (Carrollton, TX) is 0.20 miles away from my home, and has the same number of bedrooms and a similar square footage. That property was sold on 10/31/2019 for $160,000." Households could have used our proposed argument directly; however, to clarify that the content is just a suggestion, we included the following: "You can find information about this sale by searching for the property's address on Zillow.com or Redfin.com. On these websites you can find other comparable properties to support your protest." Additionally, we mentioned that subjects could protest based on different arguments, offering the following message: "You can also

[36] We identified one comparable property for all households in the subject pool, but we only displayed this information in the letter for the subjects randomly selected for the extra aid letter.

protest based on the appraised market values of comparable properties, which can be found on www.dallascad.org/SearchAddr.aspx."

We created an algorithm that identified one comparison property for each household by combining data from the tax rolls with data from recent property sales from the website Redfin.com. For each subject, the algorithm searches for properties that have been sold by late 2019 or early 2020, and were similar to the subject's own property in a number of dimensions (e.g., number of bedrooms, bathrooms, square footage built, location) but were sold for less (between 5% and 20%) than the proposed value for the subject's own property.[37] If a homeowner were to hire an agent to protest on his or her behalf, we believe it is likely that the agent would use a similar (or even the exact same) argument.

Note that our letters were not designed to eliminate the hassle costs fully. Even with the aid of our mailing intervention, households must spend time filing and keeping tabs on their protests, which has an opportunity cost (Goldszmidt et al., 2020). Additionally, some may find paperwork to be a considerably unpleasant activity (Benzarti, 2020, 2021). In that sense, our estimates will provide a lower bound of the full private costs from protesting.

There are some additional features of the letter that we summarize briefly for now but will be discussed in detail in Section 5. All letters included a table in the middle of the first page with information related to the subjects' properties' proposed values and estimated taxes. At the bottom of the first page, all letters included a URL to an online survey. To verify that the respondents were legitimate subjects and to link survey responses at the household level, we included a unique five-letter survey code for survey access. From here on, we refer to this as the *Field Survey*. The first goal of this survey was to provide a proxy for the dates that recipients opened the letters (Perez-Truglia and Cruces, 2017; Bottan and Perez-Truglia, 2020a). The survey included some questions meant to be used as outcomes in the analysis and are discussed in Section 5, below.

4.3 Subject Pool and Implementation Details

We started with the main sample of 423,607 residential single-family properties and focused on a subgroup of 78,462 of those homes, which constitutes our subject pool for the field experiment. We arrived at that subsample after applying a number of filters. For example, we excluded households that, according to the latest available data from the DCAD, had already filed a protest – because their decision had been made already, our letter could not

[37] In Appendix A.3, we provide details about this algorithm as well as some descriptive statistics.

have affected their behaviour.[38] The most important condition was to focus on households for whom our algorithm could find comparison properties that we could utilize in the extra aid message.[39] While not identical, the subject pool is quite similar to the main sample in observable, pre-treatment characteristics.[40] Additionally, Appendix A.3 shows that, consistent with successful random assignment, the observable pre-treatment characteristics are balanced across all treatment groups.

We timed the intervention so that our letters would arrive early enough to be able to affect the household's decision to protest before the upcoming deadline. We created the letters as soon as the administrative data including the 2020 proposed values became available (on May 16th, 2020). To accelerate delivery time, we used a mailing company located in Dallas County (i.e., within the same county as all of the recipients). As a result, the vast majority of the letters should have gotten to the subjects quickly. The mailing company dropped the letters off at the local post office's facility on May 20th. This company estimated that the vast majority of the letters would be delivered in the next couple of days. Consistent with this projection, we began to receive some answers to the Field Survey and visits to the website on May 21st. Moreover, the post office scans mail pieces when they reach the last mile immediately before delivery. More than 90% of the letters had been scanned by Friday, May 22nd, 2020. Based on data from previous years, most subjects file protests close to the deadline, which in 2020 was June 15th. For that reason, we feel confident that there was enough time between receipt of the letter and the protest deadline so that, for most households, the information provided in the letter could influence their decision to protest. Indeed, this view is consistent with the dates when subjects responded to the Field Survey and when they visited the project's website.[41]

4.4 Econometric Model

We want to compare the probability of protesting between subjects who were sent a letter and subjects who were not sent a letter, as well as between subjects assigned to different types of letters. The main outcome, P_i^{2020}, is an indicator variable that takes the value 100 if the subject filed a protest in the post-treatment period. We use a simple linear probability model:

[38] We initially selected a sample of 79,322 properties. However, due to a lag of a few days in the way DCAD reports the data, we had to drop 860 of them from the subject pool; with the updated data, we discovered that they had already protested by the time we mailed the letters.
[39] For more details, see Appendix A.1.2.
[40] See Appendix A.1.3.
[41] Results reported in Appendix A.4.1.

$$P_i^{2020} = \eta_0 + \eta_{basic} \cdot L_i^{basic} + \eta_{extra} \cdot L_i^{extra} + X_i^{pre} \eta_X + \epsilon_i \tag{6}$$

The variable, L_i^{basic} is an indicator that takes the value 1 if the household was mailed a basic aid letter and 0 otherwise. Likewise, L_i^{extra} is an indicator that takes the value 1 if the household was mailed an extra aid letter. Lastly, X_i^{pre} is a vector of pre-treatment controls. Unless noted otherwise, we use this exact same set of control variables in all regressions in the paper: the proposed value in levels and its annual growth, dummies for multiple owners, school and special districts, number of years since the last protest, a dummy for homestead status, the growth in the proposed value relative to the previous year, and for each year from 2015 to the previous year, a dummy indicating if the household protested in that year and the outcome of the protest. Given that this is an experiment, the only goal of using pre-treatment controls is to gain statistical power by reducing the variance of the error term (McKenzie, 2012). Additionally, we use the pre-treatment data to construct falsification tests.

4.5 Results

The regression results are presented in Table 1. All regressions are based on the same specification given in equation (6), above, but differ according to the dependent variable. The dependent variable in column (1) takes the value 100 if the owner protested directly and 0 otherwise. The basic aid letter increased the probability of protesting by 1.792 pp, an effect that is highly statistically significant (p-value<0.001). In turn, the extra aid letter increased the protest probability even more, by 3.509 pp (p-value<0.001).

We conduct a falsification test in the spirit of event-study analyses. Figure 2.b presents the results. The rightmost coefficient shows the effects of each type of letter on the probability of protesting in 2020, which are identical to the two coefficients reported in column (1) of Table 1. The rest of the coefficients correspond to the same regression specification but where the dependent variables are protest indicators for the years 2015 through 2019, instead of 2020. Since our letters had not been sent yet, they could not possibly affect the protests in prior years. As expected, the coefficients for the pre-treatment years are close to zero, statistically insignificant and precisely estimated.

One potential concern is that our letters induced protests that were ultimately not successful. We can provide direct evidence that this was not the case. In column (2) of Table 1, the dependent variable takes the value 100 if the household protested directly and won (i.e., received a discount on their market value assessment) and 0 otherwise (i.e., if either the protest was unsuccessful or if the household did not protest directly). The coefficients are still economically and statistically significant. The ratio between the coefficients on the extra aid letters from columns (1) and (2) suggests that 75% ($=\frac{2.621}{3.509}$) of the marginal protests that were induced by our letter with the extra aid were successful. This success rate is comparable to the corresponding success rate of 78% ($=\frac{6.76}{8.67}$) observed in the control group.[42] In other words, the additional protests that were induced by our letters were, on average, roughly as successful as the protests in the control group.

In Appendix A.4.2 we present a number of additional robustness checks. We show that, the increase in the direct protests induced by our letters did not crowd out the protests through agents. We report the effects of the letters on online protests separately from protests by mail. Last, we show the extent to which the marginal protests induced by our letters were successful under alternative definitions of success: the reduction in market value and the reduction in estimated taxes.

4.6 Causal Mechanisms

In this section we unpack evidence related to the causal mechanisms behind the effects of the letters.

Our favorite interpretation is that the letters increased the likelihood of protesting by mitigating the underlying hassle costs. One alternative interpretation could be that our letters simply acted as a reminder of the opportunity to protest. This explanation is unlikely in our context, however, as the proposed property taxes are quite salient around the time subjects received our letter. Moreover, we have several pieces of evidence indicating that the reminder effect was probably minor.

The first, and most direct, piece of evidence is based on a comparison between the two types of letters. If the letters acted as a simple reminder, then their effects should have been the same regardless of whether the letter included the extra aid or not. Column (1) of Table 1 indicates that the effect of the basic aid letter was positive (1.792 pp) and statistically different from the effect of the extra aid letter (3.509 pp). The difference between the coefficient estimates indicates that the extra aid message, on its own, had an effect of 1.717 pp ($= 3.509 - 1.792$), which was highly statistically significant (p-value<0.001). This constitutes demonstrates that, at most, the reminder mechanism could only explain the effects of the basic aid letter.

Moreover, we provide direct evidence that the reminder effect did not even play a significant role for the basic aid letter. For this test, we exploit heterogeneity on whether households were mailed a notification by the DCAD. Starting on May 15, all homeowners were able to download the Notice of Appraised Value by going to the DCAD webpage (we provide a sam-

[42] This success rate is based on the ratio of the share of direct protests that were successful (6.76, from column (2)) to the share of direct protests (8.67, from column (1)).

ple of this notification in Appendix E). Additionally, the DCAD sent notifications by mail to some households (e.g., whose proposed value increased relative to the previous year) but not to others. DCAD mailed the official notifications on Friday, May 15, a few days earlier than we mailed our letters, on Wednesday, May 20, so the households should have received the official notification around five days before our letter. If our basic aid letter worked primarily through a reminder effect, it should have had a larger effect on the households that did not receive the Dallas CAD notification, as the rest of the households had already been already reminded about the opportunity to protest through the official notification. The results are presented in columns (3) and (4) of Table 1, which split the sample by whether the subjects were not (column (3)) or were (column (4)) mailed a notification from the DCAD.[43] The effects of the basic aid letter on subjects that were and were not mailed a notice are on the same order of magnitude (coefficients of 1.449 and 1.935, respectively), and moreover they are not statistically distinguishable (p-value=0.317). This finding constitutes suggestive evidence that the effect of the basic aid letter went way beyond a simple reminder effect.

Regarding the effects of the basic aid message, our preferred interpretation is that they were mainly the product of the walkthroughs provided through the project's website. A first piece of evidence for this interpretation relies on a perhaps unusual source of data. In the project's website, we provided an email address to contact the researchers in case the subjects had any concerns about the research project. While this was not the purpose of providing the email address, a number of subjects sent emails to this address expressing gratitude for the letter and website and often mentioning hassle costs. For example, some subjects mentioned that they had wanted to protest for years but did not know how to until receiving our letter. Other subjects mentioned that they thought protesting was more complicated than it was until they looked at our instructions. Similarly, the Field Survey included a final, open-ended question in case the subjects wanted to share any thoughts with the researchers. Many subjects used that space to express gratitude and sometimes made explicit how the information contained in the letter and the website helped them navigate the protest process. Another source of supporting evidence for this mechanism relies on data from the traffic to the project's website. We estimate that the basic aid message generated a total of 903 additional direct protests.[44] We can compare this number of additional protests that were induced by

[43] According to the information posted on its official website, on May 15th, 2020 all Notices of Appraised Value for residential and commercial real properties became available on the Dallas CAD website. Additionally, the DCAD mailed the 2020 Notice of Appraised Value on May 15th to those properties meeting the following criteria: an increase in the property's appraised value, an ownership change, a loss of homestead exemption, rendered property, or a new property. The DCAD did not mail notices if the property value did not change or if the property value decreased. Source: `http://dallascad.org/ViewPDFs.aspx?type=1&id=%5C%5Cdcad.org%5Cweb%5Cwebdata%5Cheadlines%5CHEALTHALERTRecentHeadlines04032020.pdf`.

[44] We arrive at this figure by taking the effect of the basic aid letter (1.792 pp, from column (1)) and multiplying it by the total number of letters sent that included at least the basic aid message (50,394).

the aid message to the number of unique visits to the website. Google Analytics recorded a total of 2,769 unique visits to the walkthroughs to protest online or by mail (for more details, see Appendix A.4.1). Some of those visitors may have looked at the walkthroughs but did not protest, while some of those visitors may have used the walkthroughs but would have protested even without them. If we assume that around a third of those website visitors were induced to protest by our website, we would explain all the 903 additional protests generated by the basic aid message. In other words, it would not be far-fetched to attribute the entire effect of the basic aid message to the walkthroughs.

Regarding the extra aid message, our preferred interpretation is that subjects either used our proposed argument *as is* or followed our instructions to come up with an argument of their own. Indeed, we can provide some direct evidence that some subjects took the suggested argument in the extra aid message and used it "as is" in their protest form. This test is based on data from the 5,026 households in the subject pool who protested online. For these households, we observe the opinion of value that they entered in the online form. In column (5) of Table 1, the dependent variable takes the value 100 if the households provided an opinion of value in their protest, which is within half a percentage point of the value we selected for their extra aid message.[45] The bottom of column (5) shows that, in the control group with no letter, there was a 3.37 pp chance that a household enters an opinion of value that coincided almost exactly with the value that we would have shown them (i.e., had they been assigned to the extra aid message treatment). In other words, it is a rare coincidence for subjects to use an almost identical opinion of value as the one we would have suggested. For households that received the basic aid letter, that probability remains equally low – which is expected, because the basic aid message did not include any information about the argument. For households that received the letter with the extra aid, however, the outcome variable increases by a whooping 15.287 pp (p-value<0.001). This constitutes evidence that a substantial fraction of households that received the extra aid message copied the suggested example directly into their protest forms.

4.7 Magnitude of the Effects

A challenge with interpreting the magnitudes of the effects in mailing experiments is non-compliance: for example, some households may not have received the letters, or they may have received them but did not read them. To correct for these types of non-compliance, we need an estimate of the reading rate (i.e., the share of recipients that actually read the letter

[45] One limitation with this exercise, however, is that it is based on a subsample (households that protested online and entered a value in the opinion of value field) that is not random, and thus introduces the possibility of endogeneity bias despite the random assignment.

on time). Following Bottan and Perez-Truglia (2020a), we combine estimates from different sources to approximate the reading rate. According to the U.S. Monitor Non-Profit Standard Mail Delivery Study, around 95% of standard non-profit mailers were successfully delivered (U.S. Monitor, 2014). Based on data from the U.S. Postal Service Household Diary Survey (Mazzone and Rehman, 2019), we estimate that, conditional on delivery, around 74% of our letters were opened by the recipients.[46] If we combine the two estimates above, we arrive at a reading rate of 70.3% ($= 0.95 \cdot 0.74$). As a result, to account for this source of attenuation bias, we need to scale the coefficients up by a factor of 1.42 ($= \frac{1}{0.703}$). The resulting scaled-up effects would be 2.55 pp for the basic aid letter and 4.98 pp for the extra aid letter. This is still a conservative scale-up factor, as there may be additional sources of non-compliance.[47]

Another potential source of non-compliance is that of spillovers: i.e., if treated households shared information from the letters with neighboring households who were in the control group, that would introduce an attenuation bias in our estimates. However, we provide evidence that this form of non-compliance is negligible: Appendix A.4.3 shows that the estimated spillovers are statistically insignificant and precisely estimated at zero.

To translate the hassle costs into dollar amounts, we combine the results from the field experiment with the results on the homestead cap discussed in Section 3 above. We focus on the effect of the most complete letter (the extra aid letter). The effect of this letter gives a lower bound on the costs from protesting, as this letter did not eliminate the hassle costs completely. For example, subjects still had to follow the instructions to file the form, and may also need to take further action in the future, such as discussing a settlement in informal or formal hearings. The scaled-up effect of the extra aid letter is 4.98 pp. According to the calculations reported in section 3.2 above, each \$100 reduction in the tax amount due to the homestead cap decreases the protest probability by 2.14 pp. These results imply that the homestead cap would need to reduce the tax amount by \$232 in order to generate a reduction of 4.98 pp in the protest rate. These results imply that the average hassle cost is on the order of \$232.[48]

As a sanity check, we can compare our estimate of hassle costs to the fees charged by agents that protest on households' behalf. We identified one such company that offered the service for a flat fee. Let us assume that the marginal client of this firm is indifferent between hiring this agent or protesting on her/his own. In that case, the flat fee should constitute a

[46] This figure is based on the 2018 HDS Recruitment Sample and corresponds to the estimate of treatment of advertising mail reported in Figure 5.3 of (Mazzone and Rehman, 2019). See Bottan and Perez-Truglia (2020a) for more details.

[47] For example, some households may have opened the letter but when it was too late (i.e., after they had filed a protest or after the protest deadline, whichever came first).

[48] This is just a rough approximation. Among other things, it assumes that households only care about the costs and benefits this year, but in reality there may be dynamic considerations too.

measure of the hassle cost of the marginal client.[49] In 2020, the flat fee was \$139 for properties assessed below \$200,000, and \$305 for properties assessed between \$200,000 and \$500,000.[50] Those flat fees (\$139 and \$305) are on the same order of magnitude as our estimated average hassle cost (\$232), thus suggesting that our estimates are on the right order of magnitude.

5 Conditional Cooperation

5.1 Conceptual Framework

The above sections provide evidence suggesting that households are quite elastic to the private benefits and costs of protesting. While this evidence suggests that selfish motives are important, that does not rule out the possibility that fairness concerns matter too. More precisely, we measure conditional cooperation: i.e., whether households are more willing to tolerate taxes when they think other households in the county are paying their fair share.

Anthropologists often consider cooperation and reciprocity as examples of features that are present among all people, called "human universals" (Brown, 1991). Cooperation and fairness concerns have also been documented and studied by economists (Andreoni, 1995; Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Fehr and Schurtenberger, 2018). Multiple studies documenting conditional cooperation in the laboratory have been conducted. For instance, there is robust evidence from well-designed laboratory experiments documenting that conditional cooperation is frequent and robust when playing the public goods game, in spite of individuals' incentives to free ride and absence of apparent private benefits from cooperation (Andreoni, 1995; Gächter, 2007). In our context, conditional cooperation is rooted in similar fairness principles: households may believe it is fair to pay taxes when other households are also paying taxes. We designed a field experiment to test this hypothesis.

5.2 Design of the Field Experiment

To measure conditional cooperation, in the ideal experiment we would randomize how much others are paying in taxes and measure whether these randomly assigned average tax rates have an effect on individuals' decisions to protest. However, randomizing the average tax

[49] This is just a rough approximation. First, the marginal household may be indifferent between protesting through an agent or not protesting at all. Second, there may be additional factors at play. For example, the marginal customer may be willing to pay more than the hassle costs under the belief that that the agent can negotiate higher tax savings.

[50] Source: `https://www.dfwtaxadvisor.com/practice_areas/property-tax-protest/`, accessed on May 15th, 2020.

rate is not feasible.[51] Instead of manipulating the average tax rates directly, we manipulate the subject's perception of the average tax rate through an information-provision experiment. The goal is to leverage the fact that information on how much others pay in taxes is not easily accessible and thus probably not well-known; for example, some households may underestimate the average tax rate while others may overestimate it. Hence, by providing households with accurate information about the average tax rate, we can induce exogenous shocks to those perceptions. Then, we can measure if these information shocks affect whether the households feel their taxes are fair and whether they choose to protest them.

This information-provision experiment was the second treatment arm of the field experiment introduced in Section 4 above. This treatment arm randomized the content of the table appearing in the middle of the first page of the letter. In Figure 3, this table is highlighted inside a red box with dashed lines. This box is just for explanatory purposes and was not included in the actual letters sent to subjects. All letters included a table, but we randomized (with 50% probability) whether the table included the column with the figures for the average Dallas home (positioned in the last column). This table provided information on whether households' tax rates were above or below the average Dallas home, and by how much. Additionally, we cross-randomized whether the table included the last row, making the tax rates more salient and explicit. As stated in the RCT pre-registration, this additional treatment arm was meant for a minor robustness check and as such it is relegated to the Appendix.[52]

The random inclusion of an additional column was meant to provide a shock to households' perceptions of the tax rate other households pay. However, a naïve comparison of the protest rate among households that received the information shock and those that did not would deliver incorrect estimates to the extent that households update heterogeneously to the information shock (Cruces et al., 2013; Cullen and Perez-Truglia, 2018; Bottan and Perez-Truglia, 2020a; Fuster et al., 2018). For example, suppose that half of the sample was originally downward-biased and, thus, updated upwards, while the other half, who was originally upward-biased, updated in the opposite direction. In such a case, despite the fact that households react to the information, we would find a null average effect of information disclosure, because the two opposite effects cancel out. To overcome this limitation, and as anticipated in the pre-registration, we designed the experiment based on the disclosure-

[51] This logic is based on the implicit assumption that households care about the tax rate that other households pay. In theory, households could care about the average tax *amount* instead of the average tax *rate*. The main motivation for using the tax *rate* is that survey data indicates that the vast majority of households think that it would be fair for all households to pay the same property tax rate – results reported in Appendix A.5.1.

[52] The randomization of the last row was meant to address the concern that households compare tax rates (instead of tax amounts) because it is made salient by our letter. For more details on the design of this treatment arm, see Appendix A.3. In turn, the results from this treatment arm are reported in Appendix A.5.5. We find similar results regardless of whether the tax rate was made salient or not.

randomization design from Bottan and Perez-Truglia (2020a), which is described in detail in the next section.

The main goal is to measure the effects of the information-provision experiment on the decision to file a protest. Additionally, we included two questions from the survey included in the letter, known as the Field Survey, to be used as outcome variables. A copy of the full survey instrument is included in Appendix F. The first question simply asks individuals whether it is likely or not that they will file a protest this year. The goal of this outcome is to provide a validation exercise. The second, and most important, question was designed to provide a test of the causal mechanism at play. This question elicits whether the household thinks that their own taxes are fair or unfair relative to other households in the county, on a subjective scale from 1 to 10. The hypothesis is that, according to conditional cooperation, finding out that the average tax rate is higher should reduce the household's feeling of unfairness.

5.3 Econometric Model

Let Y_i^{post} be the outcome of interest, measured after the information provision experiment. For example, this could be an indicator of whether the individuals filed a protest post-treatment. Let τ_i be the individual's own tax rate and $\overline{\tau}$ be the *actual* average tax rate in the county. And let D_i be an indicator variable that takes the value 1 if the information on the average tax rate was shown to the subject. The regression of interest is the following:

$$Y_i^{post} = \nu_0 + \nu_1 \cdot D_i \cdot (\overline{\tau} - \tau_i) + \nu_2 \cdot (\overline{\tau} - \tau_i) + \nu_3 \cdot D_i + X_i^{pre} \nu_X + \varepsilon_i \tag{7}$$

As before, X_i^{pre} corresponds to the vector of pre-treatment control variables, which contains the same variables listed in Section 4.4 above. Again, since this is an experiment, the goal of using pre-treatment controls is to gain statistical power by reducing the variance of the error term (McKenzie, 2012). We can also use the pre-treatment data to construct falsification tests in an event-study fashion.

The coefficient ν_2 measures the relationship between the outcome and the potential information shock (i.e., $(\overline{\tau} - \tau_i)$) when the true average tax rate is not disclosed, while ν_1 measures how much stronger that relationship becomes when the true average rate is disclosed. Analogous to the design in Bottan and Perez-Truglia (2020a), the key coefficient, ν_1, measures the effects of the information shock.

A key challenge in the field experiment is that we do not observe households' prior beliefs about the average tax rate. To overcome this challenge, we take advantage of the fact that individuals systematically update their beliefs upwards or downwards depending on whether their own tax rates are below or above the average tax rate. The assumption, which we validate in the next section, is that individuals who pay more than average underestimate how much they pay relative to others and, when shown the information, they will update their perceptions upwards. On the other hand, individuals who pay less than average overestimate how much they pay, so when shown the information they will update their perceptions downwards.

5.4 Complementary Survey Experiment

To validate the design of the information-provision experiment (and the econometric model) used in the field experiment, we designed a complementary survey. This complementary survey was included in the same RCT pre-registration as the field experiment and was conducted through Amazon Mechanical Turk around the same dates as the field experiment: from June 5th to June 15th, 2020. We followed several best practices for recruiting individuals in Mturk.[53] The full survey instrument is attached as Appendix G, and is summarized next. This survey experiment starts by eliciting the (prior) beliefs about the average property tax rate in the subject's county. After that elicitation, we conduct the information-provision stage, where a random half of the subjects receive accurate information about the average tax rate in their county. Next, we re-elicit the (posterior) beliefs. With this information, we can measure how prior beliefs are distributed and how individuals update those beliefs in light of the information provided to them.

We collected responses from 2,065 U.S. homeowners.[54] The main results from the Mturk Survey are presented in Figure 5. Figure 5.a shows subjects' initial misperceptions regarding average property taxes. The x-axis corresponds to the household's actual relative taxes in 2019: that is, the difference between the average tax rate in the county and the household's own tax rate. The y-axis shows the individual's prior beliefs about the relative taxes in 2019. A slope of 1 would correspond to the case of accurate perceptions. Instead, the coefficient (0.237) falls significantly short of 1, indicating significant misperceptions. More precisely, Figure 5.a shows that the misperceptions are systematically skewed toward the middle: individuals who pay more than average underestimate how much they pay relative to others, and individuals who pay less than average overestimate how much they pay. Indeed, this type of middle-bias has been documented in a variety of settings (Cruces et al., 2013). As a result, we would expect individuals toward the left side of the x-axis in Figure 5.a to

[53] For more details about the design and implementation of this survey, see Appendix A.5.3.

[54] Appendix A.1.3 provides descriptive statistics for this sample. We show that, in terms of their observable characteristics, the Mturk sample is certainly not identical to the other samples used in this paper, but not wildly different either.

update their beliefs downwards, while individuals toward the right side of the x-axis should update their beliefs upwards.

Figure 5.b corresponds to the belief updating. This figure illustrates the intuition behind the identification strategy: we can anticipate whether the disclosure of accurate information translated into upwards or downwards revisions in beliefs, based on whether a household is toward the left or the right of the x-axis. The x-axis in Figure 5.b is the same as in 5.a, but the y-axis in 5.b corresponds to the subjects' posterior beliefs (that is, after the information-provision experiment) instead of their prior beliefs (as in 5.a). The blue dots correspond to subjects in the control group (those not shown the feedback about the true average tax rate). For this group, the relationship between perceived and real rates continues to be weak. In contrast, the red diamonds correspond to the treatment group (i.e., subjects who were shown the feedback). These red diamonds show that, as expected, the correlation between perceptions and truth becomes markedly stronger when individuals are provided with accurate feedback. This finding means that individuals who were overestimating the average tax rate updated downwards, and individuals who were underestimating it updated upwards. For a more formal test, we can compare the slope between perceptions and truth in the control group (0.154) versus the corresponding slope in the treatment group (0.613) – consistent with significant learning, the difference between the two is not only large but also highly statistically significant (p-value<0.001).

While Figure 5 shows the effects of the information shocks in a lot of detail, it is convenient to summarize those results in a single parameter. That is precisely what the econometric model from Section 5.3 was designed to do. The results from this regression specification are presented in Table 2. All columns in this table are based on the same regression specification, but use different samples and dependent variables. In column (1), the data is from the Mturk Survey respondents and the dependent variable is the posterior belief about the average tax rate in the county. *Information Shock* corresponds to the information shock in that regression specification (i.e., the term $D_i \cdot (\overline{\tau} - \tau_i)$). The coefficient on *Information Shock* from column (1) indicates that a 1 pp increase in the information shock increases the posterior belief by 0.393 percentage points. This rate of pass-through is significantly above zero, and statistically significant (p-value<0.001). To the extent that this coefficient falls short of 1, it suggests that individuals did not fully incorporate the feedback given to them, which is standard in these types of experiments and perfectly consistent with standard learning models.[55] Indeed, this coefficient of 0.393 is in the ballpark of the pass-through rates found in other survey experiments. Just to mention one example, Bottan and Perez-Truglia (2020a) uses a similar

[55] For example, in the context of Bayesian learning, individuals may not fully update their beliefs because they feel confident about the accuracy of their prior beliefs, or because they do not feel confident about the accuracy of the signal provided to them.

research design but in the context of home price expectations, and finds a pass-through rate of information shocks of 0.205.[56]

5.5 Results from the Field Experiment

We now turn to the results from the field experiment. Table 2 presents the estimation results. In column (2), the dependent variable takes the value 100 if the household protested directly in 2020 and 0 otherwise. This analysis is based on the sample of 50,394 subjects from the field experiment who were randomly selected to receive a letter. The coefficient on *Information Shock* indicates that a household that finds out that the average tax rate ($\overline{\tau}$) is 0.1 pp higher than its own tax rate decreases the probability of protesting in 2020 by 0.095 pp, which is statistically significant (p-value=0.066).

We conduct a falsification test in an event-study fashion. Figure 2.c presents the results. The rightmost coefficient shows the effect on the probability of protesting in 2020, which is identical to the coefficient on *Information Shock* reported in column (2) of Table 2. The rest of the coefficients correspond to the same regression specification but where the dependent variables are protest indicators for the years 2015 through 2019, instead of 2020. Since our letters had not been sent yet, the information shocks should not have had an effect on protests in prior years. As expected, the coefficients for the other dependent variables are close to zero, statistically insignificant and precisely estimated.

One challenge with interpreting the magnitude of the coefficient on *Information Shock* is that it is an intention-to-treat effect, due to multiple sources of non-compliance. As discussed in Section 4.7, some of the letters may not have even been opened, or may have been opened when it was too late. Additionally, even for those letters that were successfully opened, households may not have paid enough attention to the information on average taxes provided in the table. We can partially address these forms of non-compliance by focusing on the subsample of households that responded to the Field Survey. By construction, all of those households must have read the letter (otherwise they would not know the survey link and the survey code needed to fill out the survey). Furthermore, it is reasonable to assume that this subsample of survey respondents cared enough about the topic to pay close attention to the information provided in the letter.

Column (3) of Table 2 reproduces the same analysis as in column (2), except that while column (2) uses data for all the subjects that we sent letters to, column (3) is restricted to

[56] This result corresponds to the coefficient reported in column (1) of Table 2 from Bottan and Perez-Truglia (2020a).

the subsample of 1,888 of those households that responded to the Field Survey.[57] The survey respondents are certainly not a random sample. In terms of household characteristics, such as home value, number of bedrooms or tax rate, the differences between survey respondents and non-respondents are statistically significant but small.[58] However, there is one substantial difference between the samples in columns (2) and (3) of Table 2: the share of subjects who protested in 2020 (50.26%, from column (3)) is much higher relative to the corresponding share among subjects who received a letter (11.29%, from column (2)). A natural interpretation is that the subjects who paid the most attention to our letter were those who were on the fence about protesting in 2020.[59]

The coefficient on *Information Shock* is negative (-15.392) for the survey respondents (column (3) of Table 2) and statistically significant (p-value=0.006). This coefficient is much larger in magnitude than the corresponding coefficient reported in column (2). This difference is partly mechanical: since the baseline rate is much larger for survey respondents (50.26 in column (3) versus 11.29 in column (2)), it is natural for the effects to be larger too. Moreover, as discussed above, the stronger effects are probably due in great part to the fact that survey respondents were paying closer attention to the information included in the letter.

One concern with the analysis from column (3) of Table 2 is that, despite the random assignment, the endogenous nature of survey responses may introduce an endogeneity bias. To address this concern, Figure 2.d presents the event-study analysis for this specification: i.e., estimating the same regression as in column (3) of Table 2 but using as dependent variables indicator variables for whether the respondent protested in each of the years between 2015–2020. That is, Figure 2.d is identical to Figure 2.c only that it restricts the sample to survey respondents. The rationale for this exercise is that observing "effects" on the protests in pre-treatment years would suggest that restricting to the survey respondents introduced a selection bias. Reassuringly, the effects on the pre-treatment outcomes are close to zero, statistically insignificant, and precisely estimated.

Now that we are focusing on the respondents to the Field Survey, we can estimate the effects on the survey outcomes based on the questions included in that survey. The first question is about the stated intention to protest. The results are presented in column (4) of Table 2, which is identical to column (3) except that the dependent variable takes the value 100 if the household states that it is likely (or very likely) that it will protest in 2020 and takes the value zero if the household states it is unlikely (or very unlikely) to protest. Based

[57] The implied response rate to the survey, 3.7%, may seem low at first glance, but it is substantially higher than the response rates in comparable studies that sent a survey link through letters (e.g. Perez-Truglia and Troiano, 2018; Bottan and Perez-Truglia, 2020a).

[58] Results reported in Appendix A.5.6.

[59] Additionally, it is possible that households that found our letter helpful to protest wanted to reciprocate our help by responding to our survey.

on the evidence that the information shock affected the actual decision to protest (column (3)), we would expect to see similar effects on the intention to protest (column (4)). This is exactly what we find: the coefficient on *Information Shock* from column (4) is negative (-13.220), highly statistically significant (p-value=0.008), and close in magnitude to (and statistically indistinguishable from) the corresponding coefficient from column (3) (-13.220 in column (4) vs. -15.392 in column (3)).

In Appendix A.5.5 we present a number of additional robustness checks. For example, we show that the effects of the information shock are similar regardless of whether the tax rate was made salient or not. Additionally, we show that the effect of the information shock is consequential not only for the number of protests but also for the subsequent market values and tax amounts.

5.6 Causal Mechanisms

The evidence presented above suggests that the information about the average tax rate had a significant effect on the decision to protest. Next, we discuss, and provide evidence on, some of the potential mechanisms at play.

Our favorite interpretation is that households changed their perception of the average tax rate and that lowered their feelings of unfairness. To probe this mechanism, we leverage the other question that we included in the Field Survey, on the household's perceived unfairness of its property taxes. If this mechanism is at play, we would expect that an increase in the perceived average tax rate would reduce how unfair the household views its own taxes. The results are presented in column (5) of Table 2, which is identical to column (3) except that the dependent variable measures the household's feeling of unfairness in a scale from 0 (very fair) to 10 (very unfair). As expected, the coefficient on *Information Shock* from column (5) is negative (-0.468) and statistically significant (p-value=0.060). This effect is large in magnitude too, at least when compared to the effect on the protest choice. For example, an information shock of 0.1 pp causes an increase of 2.14% of the standard deviation in unfairness (column (5)), while it causes an increase of 3.08% of the standard deviation in protests (column (3)).

In addition, in the MTurk Survey we included some similar questions to assess the role of fairness: we asked subjects whether their taxes are unfair relative to other households in the county, whether their taxes are too low or too high, and we elicited the tax rate that the household would consider the most fair for it to pay (holding constant the tax rates of everyone else). Consistent with the effects reported above in the field experiment, we find that the information shocks from the Mturk Survey affected the fairness outcomes as well (results reported in Appendix A.5.4).

While conditional cooperation is our preferred interpretation of the effects of the information shocks, we discuss an alternative mechanism. In theory, it is possible that subjects reacted to the information on the average tax rate because they learn from that information about whether they will be successful if they protest. More precisely, a subject who finds out that it is paying a higher tax rate than the average household in the county may infer from that information that the reason for that difference is that other households have protested in the past and were successful. However, below we provide evidence against this alternative mechanism.

First of all, if we assume that households process the information rationally, we should rule out this mechanism from the very beginning. This is because, by construction, a successful protest has a negligible effect on the tax rate that the household pays. Specifically, the tax rate is computed by dividing the tax amount by the proposed value of the property. When a household has a successful protest, it reduces the value of both the numerator and the denominator, thus leaving the tax rate approximately unchanged. As a result, when a household finds out that the average tax rate is 1 pp higher than its own, it would be irrational for that household to infer anything about whether other households were successful at protesting or not (for a more detailed discussion, see Appendix A.5.2).

It is still possible, however, that households make inferences about the protest process from the information on the average tax rate because they are processing the information irrationally. To address that remaining concern, we provide a test that leverages the fact that this alternative mechanism should affect households in an asymmetric fashion. In column (6) of Table 2, we include the households that would be most affected by this alternative mechanism: those who had not protested in past years and who find out that their own tax rates are above the county average. These households may (incorrectly) infer that the reason why other households are paying lower tax rates than them is because those other households protested and were successful but they did not protest themselves. In column (7), we include the rest of the households (i.e., those that either protested in the past or find out that they pay below-average tax rates), for whom this alternative mechanism would not apply, or at the very least not as strongly. We do not find any evidence of the type of asymmetry predicted by the alternative mechanism: the coefficients from columns (6) and (7) are similar in magnitude (-11.05 vs. -14.88) and their difference is statistically insignificant (p-value=0.59). Indeed, if anything, the difference in point estimates goes in the opposite direction as the one predicted by the alternative mechanism. In sum, while we cannot rule out that the alternative mechanism plays some role, it is not plausible that it fully explains the effects of the information shocks reported above.

5.7 Magnitude of the Effects

To illustrate the magnitude of conditional cooperation, we need to account for non-compliance. The first form of non-compliance is that the subjects may have not read the letter. As explained above, we can address this form of non-compliance by focusing on the results from column (3) of Table 2, because for this sample we are confident that the subjects read the letter. However, a second form of non-compliance still remains as explained above: even if they read the letter, households are not expected to fully incorporate the feedback into their beliefs. We can use the results from the Mturk Survey to correct for this additional form of non-compliance. Because each additional 1 pp in the information shock increased the perceived average tax rate by 0.393 pp (column (1) of Table 2), we can use a scale-up factor of 2.54 ($= \frac{1}{0.393}$). Scaling-up the coefficient on *Information Shock* from column (3) of Table 2 implies that increasing a household's perception of the average tax rate paid in the county by 0.1 pp would decrease the protest probability by 3.90 pp ($= 0.1 \cdot 15.392 \cdot 2.54$). The 0.1 pp reduction in the average tax rate corresponds to a 4.7% reduction relative to the baseline tax rate in this sample (2.11%), while the 3.90 pp reduction in the protest probability corresponds to 7.7% of the average protest probability reported in column (3) of Table 2. These two estimates suggest that the elasticity of the protest probability with respect to the perceived average tax rate is 1.63 ($= \frac{7.7}{4.7}$); in other words, households are elastic to their perceptions of the taxes paid by everyone else, yet not nearly as elastic as they are to the private benefits (i.e., the 11.26 elasticity reported in the analysis of the homestead cap of Section 3).

6 Heterogeneity by Partisanship

Survey data indicates that Republicans and Democrats *say* that they have different opinions on taxation and redistribution, with Republicans typically expressing preferences for lower taxes and lower income redistribution (Di Tella et al., 2017; Alesina et al., 2020; Stantcheva, 2020). Motivated by this fact, in this section we assess whether those partisan differences materialize in the decision to protest taxes.

6.1 Data on Party Affiliation

To split the analysis by Republican and Democrat households, we use the information on homeowners' full names and addresses to merge, at the individual-level, the taxpayer records with the voter files. In Texas, households do not have to report a political party when registering to vote. However, whether they voted in a primary election is a matter of public record. As a result, participation in primaries provides a natural measure of party affiliation: e.g., an individual who participated in Democratic primaries, but not in Republican primaries, would be classified as Democrat. Moreover, we obtained the voter file records from a private vendor (Aristotle International) that supplements the data on partisanship from the voter files with data from other sources. For example, if an individual contributed over $200 to a Democrat or Republican candidate it is a matter of public record (Perez-Truglia and Cruces, 2017) and the vendor uses that data to infer political affiliation. The proxy for political party provided by the vendor is highly consistent with voting data: at the precinct-level, there is a 0.78 correlation between the proxy for party affiliation and the actual share of votes in the 2012 presidential election.[60]

We classified each of the 423,607 single-family homes in the main sample as more likely to identify as Republican or Democrat.[61] More specifically, we identified 57% of these households as more likely to be Democrats and the remaining 43% as more likely to be Republican. Indeed, this narrow lead in support for the Democrat party is consistent with recent electoral results from Dallas County: for example, in the 2012 presidential election, Barack Obama received 57% of the votes while Mitt Romney received 42% of the votes (and the remaining 1% of the votes went to third-party candidates).

6.2 Protest Rates by Party Affiliation

We begin by looking at differences in the rate of protesting between Democrats and Republicans. Given that in survey data Republicans are substantially more likely to oppose taxes and income redistribution (Di Tella et al., 2017; Alesina et al., 2020; Stantcheva, 2020), we would expect that Republicans are substantially more likely to protest their property taxes than Democrats. Instead, we find that, on average, Republicans are just a bit more likely to protest than Democrats: in the sample of 423,607 households, 9.05% of Republicans protested directly in 2020 while 7.92% of Democrats protested directly. Moreover, those small differences arise in great part due to differences in wealth. The top half of Figure 6.a shows the protest probabilities for Republicans and Democrats by groups of different market values. For example, the first group corresponds to properties valued under $100,000, and the last group corresponds to properties valued above $500,000. The bottom half of the figure shows the percent of Democratic and Republican homes that fall into each category. Relative to Democrats, Republicans tend to live in more expensive homes. Since owners of

[60] Results presented in Appendix A.6.1.

[61] Some households have missing information, for example because they are not registered to vote. Thus, we use a simple predictive model to impute their party affiliation. For details on this imputation, see Appendix A.6.1.

more expensive homes are more likely to protest, that could mechanically explain the differences in protest rates across partisan lines. Figure 6.a shows that, when comparing homes of roughly similar value, the differences in protest rates between Democrats and Republicans are even smaller.[62] For example, for homes in the modal category, valued between \$100K–\$200K, the protest probabilities are 6.34 pp for Republicans versus 5.82 pp for Democrats – this difference is statistically significant (p-value=<0.001) but small in magnitude. For more expensive properties, the relation flips, with Democrat households being a bit more likely to protest than Republicans in comparable homes.

6.3 Motives for Protesting

Even if Democrats and Republicans protest with a similar probability, it is possible that they protest for different motives. In Sections 3–5, we analyzed the role of different motives. Next, we break that analysis down by political party.

We start with the response to private benefits. Figure 7 reproduces the results from Figures 1.a and 1.c (on the effects of the homestead cap) but broken down by Democrats (left panels: (a) and (c)) versus Republicans (right panels: (b) and (d)) households. Being \$10,000 above the homestead cap threshold causes a decrease in the protest rate of 4.25 pp (p-value<0.001) among Democrats (from Figure 7.c) and a corresponding decrease of 4.70 pp (p-value<0.001) among Republicans (from Figure 7.d). These effects are close to each other and statistically indistinguishable (p-value=0.597), thus suggesting that the responsiveness to private benefits seem quite similar between Republicans and Democrats.

The differences are still small if we normalize the effects on protest rates by the corresponding effects on the tax amounts (Figures 7.a and 7.b). Being \$10,000 above the homestead cap threshold causes a decrease in the tax amount by \$220 (p-value<0.001) among Democrats (from Figure 7.a) and a corresponding decrease of \$204 (p-value<0.001) among Republicans (from Figure 7.b), with the difference between these two coefficients being small and statistically insignificant (p-value=0.674). In other words, for Democrats a \$100 reduction in the tax amount due to the homestead cap is associated to a 2.03 pp drop in the protest rate, while the corresponding effect is 2.31 pp for Republicans. This evidence suggests that both Republicans and Democrats are highly responsive to the private benefits from protesting, but Republicans seem, if anything, to be slightly more responsive.

Next, we look at the partisan differences in responses to private costs. These results are presented in the last two columns of Table 1. These columns reproduce the results from the baseline specification of column (1) of Table 1, but are broken down into the subsamples of

[62] The results are similar if we include protests through agents – see Appendix A.6.2.

Democrats (column (6)) and Republicans (column (7)). As with the response to the private benefits, we find the response to the private costs to be qualitatively consistent between Democrats and Republicans: the coefficients on the basic and extra aid letters are positive and highly statistically significant in both columns (6) and (7). Quantitatively, the differences are mixed. On one hand, the coefficient on basic aid letter is higher for Democrats (1.943, from column (6)) than for Republicans (1.509, from column (7)), but this difference is small and statistically insignificant (p-value=0.391). On the other hand, the coefficient on the extra aid letter is higher for Republicans (3.994, from column (7)) than for Democrats (3.027, from column (7)), with the difference being statistically significant (p-value=0.065). Taken jointly, the evidence suggests that both Republicans and Democrats are highly elastic to the private costs, but Republicans may be somewhat more elastic.

The evidence on the role of private costs and private benefits so far suggests that both Republicans and Democrats protest due to selfish motives, although Republicans may be somewhat more responsive to selfish motives. Next, we turn to the partisan heterogeneity in conditional cooperation. While fairness concerns are universal, results from laboratory experiments suggest that there are large differences across individuals in the strength of conditional cooperation.[63] Based on that evidence, it is at least possible that Democrats and Republicans differ in the strength of conditional cooperation. The results are presented in the last two columns of Table 2, which break down the baseline results from column (3) by Democrat households (column (8)) and Republican households (column (9)). The coefficient on *Information Shock* for Democrats (-18.317, from column (8)) is larger in absolute value than the corresponding coefficient for Republicans (-10.922, from column (9)). While these point estimates suggest that conditional cooperation is somewhat stronger among Democrats than among Republicans, we have to take this result with a grain of salt because their difference is statistically insignificant (p-value=0.515). In sum, the evidence suggests that conditional cooperation exists for both Democrats and Republicans, but is perhaps somewhat stronger among Democrats.

6.4 Discussion

If Republican and Democrats are so different in their views about redistribution, why do they behave so similarly when it comes to protesting their property taxes? One simple explanation for this finding is that survey data does not provide a trustworthy measure of the true preferences around taxation and redistribution. Democrats and Republicans may, for example, be more inclined to "state" different attitudes on these topics as cheap talk, to

[63] For example, while some subjects are willing to match one-to-one the contributions made by others, others prefer to match partly, and others do not care about the contributions of others at all (Gächter, 2007).

align with the party preferences.[64] When there is real money at stake, however, our results suggest that those differences disappear. To the extent that Republicans and Democrats state different views about taxation but end up behaving similarly, our evidence provides support for the aphorism that goes "everyone's a Republican on tax day".

However, not all survey questions are created equal. We provide some evidence that survey data may exaggerate the differences between Democrats and Republicans but only when asked in vague terms. Let us start with Figure 6.b, which shows the distribution of one of the most widely-used measures of preferences for redistribution: whether the government should reduce the income differences between the rich and the poor on a scale from 1 (should not) to 7 (should).[65] Using responses from the 2006-2018 waves of the General Social Survey, we break down the preferences for redistribution by political party. These survey responses suggest that Republicans and Democrats are worlds apart in their support for redistribution. For instance, while 9.5% of Republicans respond the highest redistribution category, the proportion of Democrats reporting this answer over three times as high (29.2%).

In turn, Figure 6.c is based on a different question from the General Social Survey that, instead of using a vague question, it elicits support for a specific tax: "Do you consider the amount of federal income tax which you have to pay as too high, about right, or too low?" Figure 6.c shows that the share of Democrats who think federal taxes are too high (60.5%) is a bit lower, but just slightly, than the corresponding share of Republicans (62.6%). In other words, when asked in more concrete terms, the differences between Democrats and Republicans are smaller and more consistent with the revealed-preference data from our study. Indeed, we find consistent results using other survey questions. For instance, Figure 6.d is based on a question that we included in the Mturk Survey and is identical to the question from Figure 6.c but asks about property taxes instead of the federal income taxes: "Do you consider the amount of property taxes you pay to be too low, about right, or too high?" The results from Figure 6.d are consistent with the results from Figure 6.c: the share of Democrats saying that property taxes are too high (36.9%) is a bit lower, but not much lower, than the corresponding share of Republicans (42.9%).

In Appendix A.6.3 we show that the results are similar for other survey questions on specific tax policies or that make explicit a requirement to balance government revenues and expenditures. Specifically, the results show relatively low levels of political polarization in the responses to a question that makes explicit that higher taxes are required for the provision of high level of government services and to a question on social norms about the fair distribution of a given a amount of taxes between two properties of different values.

[64] Indeed, there is evidence that in a number of topics such as estate taxation there are specific words that may cue partisan responses (Gentzkow and Shapiro, 2010).

[65] For more details about this question, see the notes to Figure 6.

Indeed, these findings are related to evidence from Stantcheva (2020) that, in the context of the effects of tax policies, there is more or less partisan polarization depending on the question used: e.g., there is more polarization when asked about the broad effects of the tax policies (e.g., the overall effect on the economy) than when they are asked about specific effects of tax policies (e.g., the effect on individual taxpayers).

7 Conclusions

The choice to file a protest of property taxes provides a unique opportunity to study support for taxation and redistribution via revealed preferences. Using experimental and quasi-experimental methods, we provided evidence on the determinants of the decision to protest property taxes. We showed that households are highly elastic to their private benefits and private costs from protesting, which suggests that selfish motives are important. We showed that fairness concerns are also a significant motive: consistent with conditional cooperation, households are more willing to pay a higher tax rate when they perceive that the average household faces a higher tax rate, too. We documented some partisan differences in the decision to protest.

While in this study we used our field experiment to study what motivates tax protests and support for taxation, we believe that this same framework can be adapted to study other questions from fields such as political economy, public economics, finance, and behavioral economics. Our setting has a number of features that we believe can make it attractive to researchers. The effects on behavior are measured with objective data from administrative records in a naturally occurring context and are based on high-stakes choices. The experiment can be conducted entirely based on publicly-available data, without the need of non-disclosure or other agreements. The experiment can be implemented in a few weeks and the final results may be ready in a couple of months. The mailing experiment is relatively cheap, costing less than $0.25 per subject. Lastly, the experiment can be implemented on massive scales; in Dallas County alone, it can potentially involve hundreds of thousands of subjects at a time and be scalable to millions of subjects by pooling multiple counties. In this spirit, we provide detailed accounts of the implementation and the data sources that other researchers can follow. Moreover, we are happy to share data, code, tips or any additional resources.

To evaluate the generalizability of the results of our field experiment, we follow the SANS conditions (Selection, Attrition, Naturalness, and Scalability) in List (2020). In terms of selection, the subject pool is highly representative of the target population.[66] This is an experimental context in which attrition is not an issue because the outcomes can be tracked

[66] For details, see Appendix A.1.3.

with administrative data. In terms of naturalness of the setting, we conduct a natural field experiment that is ideal in this dimension (Harrison and List, 2004). The natural, high-stakes setting is particularly important given the focus on social preferences (Levitt and List, 2007; Al-Ubaydli and List, 2013). Due to the type of intervention (a mailing experiment) it would be straightforward to scale it to the universe of households – moreover, there is no obvious reason to expect a "voltage drop" with the scaling up (Al-Ubaydli et al., 2017).

In the terminology introduced by List (2020), we view our study as a Wave 1 study, focusing on establishing initial causality and producing first tests of theory. Although our evidence is based on data from a specific U.S. county, to the extent that tax protests work similarly across counties in and out of Texas, the results should be generalizable to those other settings. Moreover, our field experiment and quasi-experiments could be replicated in other counties too: the homestead cap for property taxes exists in many other counties, and mailing interventions as ours should be feasible in all counties. Last, we hope our research design will be applied to other settings, to provide more evidence on the underlying causal mechanisms and other mediating factors.

Finally, while the main focus of this paper is to study tax protest motivations via revealed preferences, some of the findings may serve as key inputs for the design of the protest process. For instance, because the private costs and benefits are so important in the decision to protest, the current protest system may create undesired inequities. Our evidence supports the view that richer households protest more just because they stand to gain more from doing so. Similarly, our experimental evidence suggests that some groups – such as racial and ethnic minorities – may have unequal access to protests if they find it more difficult to navigate the protest process. Indeed, our experimental findings could explain the large differences in access to tax protests documented with observational data (Avenancio-León and Howard, 2019).

More importantly, our experimental findings hint at some low-cost interventions (nudges) that can alleviate inequities in the system (Thaler and Sunstein, 2009). Because households are so responsive to aid, access to protests may be made more equal by targeting aid to disadvantaged groups (Finkelstein and Notowidigdo, 2019). For example, holding the tax amount constant, Hispanic households were 3.61 pp less likely to protest than White households.[67] In comparison, we find that our letter with extra aid increased the protest rate of Hispanic households by 2.55 pp (p-value$<$0.001).[68] Thus, low-cost interventions of this type targeted at disadvantaged populations – promoted internally by county assessors' offices or externally through NGO's – could go a long way in making the system to protest taxes more equitable.

[67] For details, see Appendix A.1.4.

[68] This result corresponds to the regression from column (1) of Table 1, but estimated with the 20.28% of the sample that was classified as Hispanic.

References

Al-Ubaydli, O. and J. A. List (2013). On the Generalizability of Experimental Results in Economics. *In Frechette, G. & Schotter, A., Methods of Modern Experimental Economics, Oxford University Press.*

Al-Ubaydli, O., J. A. List, D. LoRe, and D. Suskind (2017). Scaling for Economists: Lessons from the Non-Adherence Problem in the Medical Literature. *Journal of Economic Perspectives 31*(4), 125–144.

Alesina, A. and G. M. Angeletos (2005). Fairness and redistribution: US vs Europe. *American Economic Review 95*, 913–935.

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

Alesina, A. and E. Glaeser (2004). *Fighting Poverty in the US and Europe.* Oxford University Press.

Alesina, A., A. Miano, and S. Stantcheva (2020). The Polarization of Reality. *AEA Papers and Proceedings 110*, 324–328.

Andreoni, J. (1995). Cooperation in Public-Goods Experiments: Kindness or Confusion? *American Economic Review 85*(4), 891–904.

Avenancio-León, C. and T. Howard (2019). The Assessment Gap: Racial Inequalities in Property Taxation. *SSRN Working Paper No. 3465010*.

Benabou, R. and J. Tirole (2006). Belief in a just world and redistributive politics. *Quarterly Journal of Economics 121*(2), 699–746.

Benzarti, Y. (2020). How Taxing is Tax Filing? Using Revealed Preferences to Estimate Compliance Costs. *American Economic Journal: Economic Policy 12*(4), 38–57.

Benzarti, Y. (2021). Estimating the Costs of Filing Tax Returns and the Potential Savings from Policies Aimed at Reducing These Costs. *Tax Policy and the Economy 35*.

Bhargava, S. and D. Manoli (2015). Psychological Frictions and the Incomplete Take-Up of Social Benefits: Evidence from an IRS Field Experiment. *American Economic Review 105*(11), 3489–3529.

Bolton, G. E. and A. Ockenfels (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review 90*(1), 166–193.

Bottan, N. and R. Perez-Truglia (2015). Losing my religion: The effects of religious scandals on religious participation and charitable giving. *Journal of Public Economics 129*.

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

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

Brown, D. (1991). *Human Universals Hardcover*. Philadelphia, PA: Temple University Press.

Cabral, M. and C. Hoxby (2012). The Hated Property Tax: Salience, Tax Rates, and Tax Revolts. *NBER Working Paper No. 18514*.

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.

Chetty, R. (2015). Behavioral Economics and Public Policy: A Pragmatic Perspective. *American Economic Review 105*(5), 1–33.

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

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

Cullen, Z. and R. Perez-Truglia (2018). How Much Does Your Boss Make? The Effects of Salary Comparisons. *NBER Working Paper No. 24841*.

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*.

Di Tella, R., J. Dubra, and A. L. Lagomarsino (2017). Meet the Oligarchs: Business Legitimacy, State Capacity and Taxation. *NBER Working Paper No. 22934*.

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.

Epper, T., E. Fehr, and J. Senn (2020). Other-regarding preferences and redistributive politics. *University of Zurich Working Paper No. 339*.

Fehr, E. and K. M. Schmidt (1999). A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics 114*(3), 817–868.

Fehr, E. and I. Schurtenberger (2018). Normative foundations of human cooperation. *Nature Human Behaviour 2*(7), 458–468.

Finkelstein, A. and M. J. Notowidigdo (2019). Take-up and targeting: Experimental evidence from SNAP. *The Quarterly Journal of Economics 134*(3), 1505–1556.

Fisman, R., P. Jakiela, and S. Kariv (2017). Distributional preferences and political behavior. *Journal of Public Economics 155*, 1–10.

Fisman, R., S. Kariv, and D. Markovits (2007). Individual Preferences for Giving. *American Economic Review 97*(5), 1858–1876.

Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2018). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *NBER Working Paper No. 24767*.

Gächter, S. (2007). Conditional cooperation: Behavioral regularities from the lab and the field and their policy implications. In *Economics and psychology: A promising new cross-disciplinary field.* Cambridge, MA, US: MIT Press.

Gentzkow, M. and J. M. Shapiro (2010). What Drives Media Slant? Evidence From U.S. Daily Newspapers. *Econometrica 78*(1), 35–71.

Goldszmidt, A., J. A. List, R. D. Metcalfe, I. Muir, V. K. Smith, and J. Wang (2020). The Value of Time in the United States: Estimates from Nationwide Natural Field Experiments. *NBER Working Paper No. 28208*.

Goolsbee, A. (2006). The Simple Return: Reducing America's Tax Burden Through Return-Free Filing. *The Hamilton Project Discussion Paper No. 2006-04*.

Hallsworth, M., J. A. List, R. D. Metcalfe, and I. Vlaev (2017). The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. *Journal of Public Economics 148*, 14–31.

Harrison, G. W. and J. A. List (2004). Field Experiments. *Journal of Economic Literature 42*(4), 1009–1055.

Holz, J. E., J. A. List, A. Zentner, M. Cardoza, and J. Zentner (2020). The $100 Million Nudge: Increasing Tax Compliance of Businesses and the Self-Employed using a Natural Field Experiment. *NBER Working Paper No. 27666*.

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

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

Kuziemko, I., M. I. Norton, E. Saez, and S. Stantcheva (2015). How elastic are preferences for redistribution? Evidence from randomized survey experiments. *American Economic Review 105*, 1478–1508.

Levitt, S. D. and J. A. List (2007). What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World? *Journal of Economic Perspectives 21*(2), 153–174.

List, J. A. (2020). Non est Disputandum de Generalizability? A Glimpse into The External Validity Trial. *NBER Working Paper No. 27535*.

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

Mazzone, J. and S. Rehman (2019). The Household Diary Study Mail Use and Attitudes in FY 2018. *Retrieved March 28, 2020, from https://www.prc.gov/dockets/document/109368*.

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

Meltzer, A. and S. Richard (1981). A rational theory of the size of government. *Journal of Political Economy 89*(5), 914–927.

Perez-Truglia, R. and G. Cruces (2017). Partisan interactions: Evidence from a field experiment in the United States. *Journal of Political Economy 125*(4), 1208–1243.

Perez-Truglia, R. and U. Troiano (2018). Shaming Tax Delinquents. *Journal of Public Economics 167*, 120–137.

Pomeranz, D. (2015). No Taxation Without Information: Deterrence and Self-Enforcement in the Value Added Tax. *The American Economic Review 105*(8), 2539–2569.

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

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

Sunstein, C. R. (2021). *Sludge.* Cambridge, MA: MIT Press.

Thaler, R. H. and C. R. Sunstein (2009). *Nudge: Improving decisions about health, wealth, and happiness.* Penguin, 2009.

U.S. Monitor (2014). 7 Myths of Direct Mailing. *Retrieved March 28, 2020, from https://www.targetmarketingmag.com/promo/7MythsofDM.pdf*.

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.

Figure 1: Effects of the Homestead Cap on the Tax Rate and on the Probability of Protesting

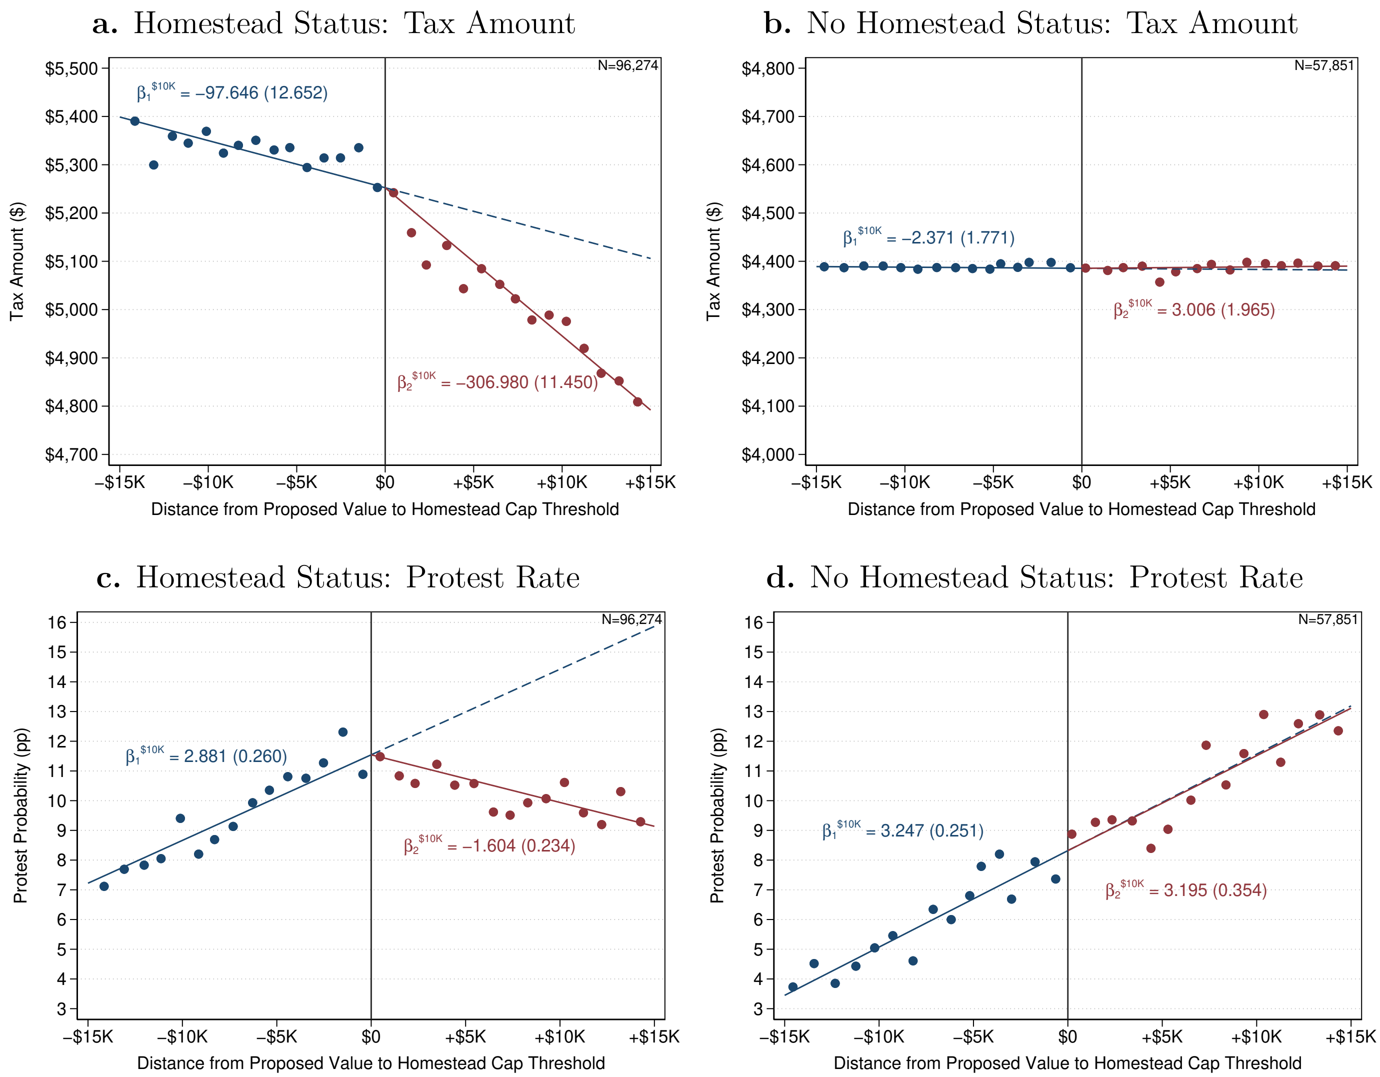

Notes: This figure features binned scatterplots of the relationship between a given outcome (indicated on the y-axis of each panel) and the distance between the 2020 proposed value and the 2020 potential homestead cap threshold (defined as 110% of the appraised value in the previous year). All regressions control for the proposed value, a dummy for whether the household protested in the previous year, and a set of school district dummies. The sample is restricted to properties for which the proposed value is within $15,000 of the potential homestead threshold. The lines correspond to linear regressions, with normalized slopes reported next to them along with robust standard errors (in parentheses) and the number of households (in brackets in the top right corner). The panels on the left half ((a) and (c)) correspond to households with 2020 homestead status, while the panels on the right half ((b) and (d)) correspond to households without 2020 homestead status. The dependent variables are: *Tax Amount* is the estimated tax amount based on 2020 proposed values and P_{2020} is an indicator variable that takes the value 100 if the household protested directly in 2020 and 0 otherwise.

Figure 2: Event-Study Falsification Tests

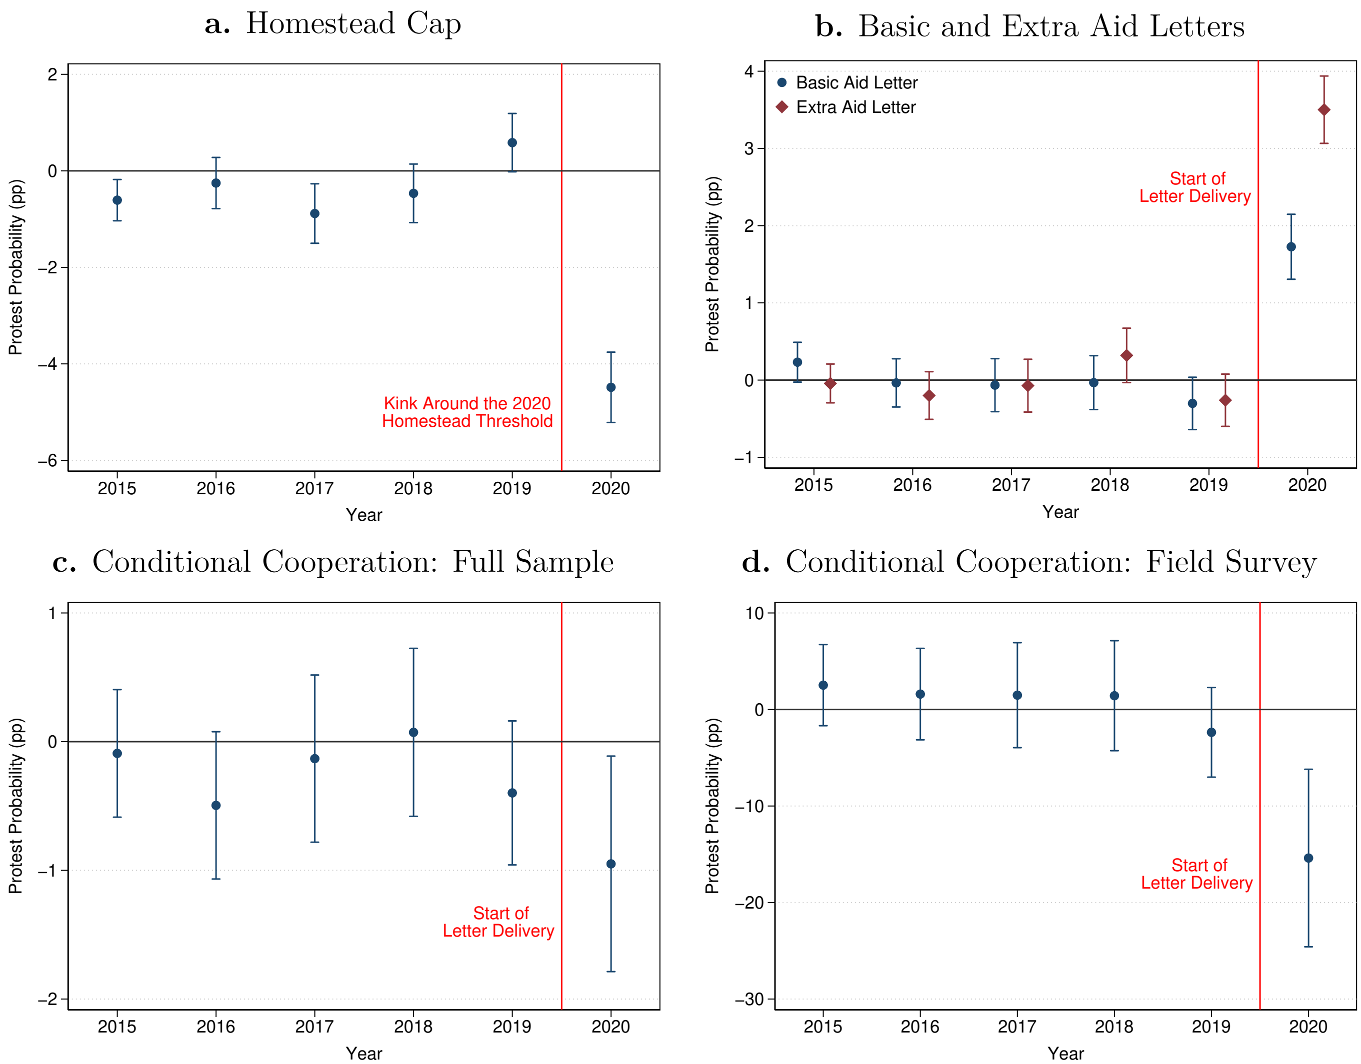

Notes: Point estimates with 90% confidence intervals in brackets, based on robust standard errors. The point estimates are computed in the same way within each of the four panels: the point estimates within each panel only change the focal year. Panel (a): the blue dots represent the difference between the slopes before and after the threshold as in Figure 1.c, but varying the year. Results based on single-family homes with 2020 homestead status and an absolute difference between the proposed value and the potential homestead cap of less than \$15,000. Panel (b): The blue dots represent the effects of the basic aid letter (relative to the no letter group), while to red diamonds represent the effect of the extra aid letter. Panel (c): The blue dots represent the coefficient on the information shock ($D_i \cdot (\overline{\tau} - \tau_i)$) based on equation (7) from Section 5.3. Panel (d): same as panel (c) except that it is based on the subsample of 1,888 subjects who responded to the Field Survey.

Figure 3: First Page of the Sample Letter

THE UNIVERSITY OF TEXAS AT DALLAS
Naveen Jindal School of Management

May 15th, 2020

Dear Joan Robinson,

We are researchers at The University of Texas at Dallas and we are reaching out to you as part of a research study**. You can lower your tax burden by protesting the taxable value assessment of your property.** We want to share information that we hope will be useful**.**

Some people may choose to protest because they feel they are paying more than their fair share**.** Find below some information about the estimated 2020 taxes for your home at 5329 Jordan Ridge D (Dallas, TX) in Dallas County*:*

	YOUR HOME	**AVERAGE DALLAS HOME**
Proposed Value	*$174,810*	*$294,846*
Estimated Tax Amount	*$3,057*	*$5,916*
Estimated Tax Rate	*1.75%*	*2.01%*

Source: Data provided by Dallas Central Appraisal District (CAD). Proposed Value is Dallas CAD's estimate of the home's market value as of January 1st, 2020. Estimated Tax Amount is our estimate of taxes due this year using the latest tax rates available (some exemptions might not be included). Estimated Tax Rate is the estimated tax amount divided by Proposed Value. Average Dallas Home values are based on all single-family homes in Dallas County, excluding condos, townhomes, and mobile homes.

The deadline to protest is June 15th, 2020**.** You can fill out a short form online or mail it in**.** You can find instructions on how to do this on the study's website:

https://www.utdallas.edu/taxproject/

If you would like to help us with our study, we kindly ask you fill out the following confidential survey**.** It only takes a couple of minutes, and we would greatly appreciate your participation:

Visit http://www.utdallas.edu/taxsurvey/ and enter validation code **AAFOGD**

800 W. Campbell Road
Richardson, TX 75080

Website: https://www.utdallas.edu/taxproject/

Please recycle

43137

Notes: A sample of the first page of the letter used in the field experiment. The information in the table varied by treatment group. Sample tables for every treatment group are presented in Figure A.4. The table appears inside a red frame with dashed lines (this frame was added to this figure for emphasis but does not appear in the actual letters).

Figure 4: Second Page of the Sample Letter

a. Extra Aid Message: No

Your household was randomly chosen to receive this letter. *We will not send you any more letters in the future.* If you have any questions about the study, you can find contact information on the study's website.

Thank you for your attention!

Alejandro Zentner
Associate Professor
University of Texas at Dallas

43137
JOAN ROBINSON
5329 JORDAN RIDGE DR
DALLAS, TX 75236-1895

b. Extra Aid Message: Yes

If you'd like to file a protest, it is really simple. You do not need an agent. You do not need to attend a hearing if you accept an online settlement offered by the county. If the county schedules a hearing and you do not attend it, the protest will simply be dismissed with no penalty.

When you protest you need to provide an argument in a few sentences. For example, you may argue that the appraised market value is too high. In that case, you could use the following:

☑ Value is over market value   Opinion of value: *$160,000*

And remember to attach a separate page (or file, if protesting online) with your argument:

I found a home that is similar to mine but was recently sold for less than my home's appraised market value. The property located at 5148 Ronryan Rd (Dallas, TX) is 0.29 miles away from my home, and has the same number of bedrooms and a similar square footage. That property was sold on 10/31/2019 for $160,000.

You can find information about this sale by searching for the property's address on Zillow.com or Redfin.com. On these websites you can find other comparable properties to support your protest. You can also protest based on the appraised market values of comparable properties, which can be found on www.dallascad.org/SearchAddr.aspx.

Your household was randomly chosen to receive this letter. *We will not send you any more letters in the future.* If you have any questions about the study, you can find contact information on the study's website.

Thank you for your attention!

Alejandro Zentner
Associate Professor
University of Texas at Dallas

43137
JOAN ROBINSON
5329 JORDAN RIDGE DR
DALLAS, TX 75236-1895

Notes: A sample of the second page of the letter used in the field experiment. Panel (a) does not contain the Extra Aid Message, while panel (b) does in the section framed with the red dashed lines. This red frame in panel (b) was added to this figure for emphasis but does not appear in the actual letters.

Figure 5: Prior Misperceptions and Belief Updating in the Mturk Survey

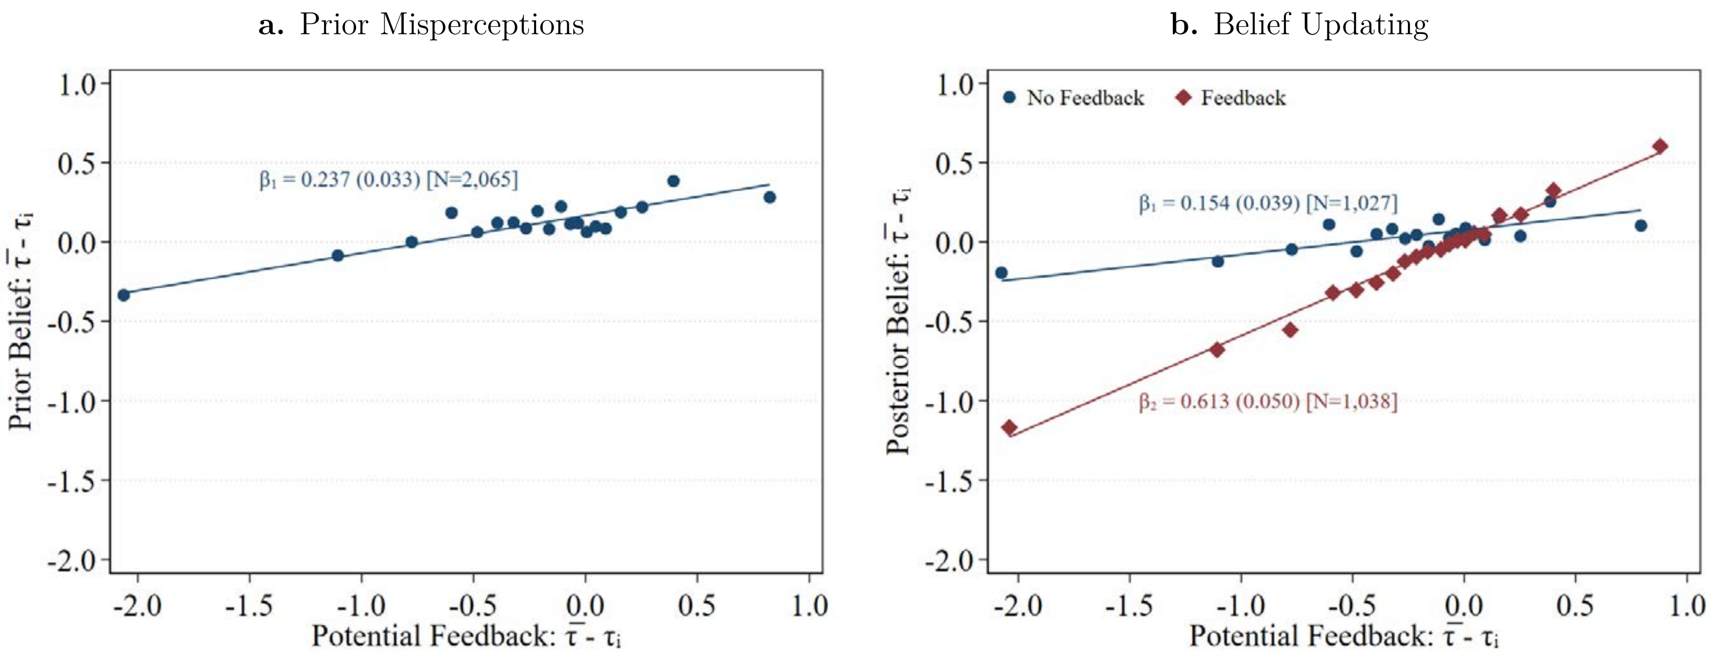

Notes: This figure shows binned scatterplots based on the Mturk Survey. Each line corresponds to a separate OLS regression, with robust standard errors in parentheses and the number of observations in brackets. $\overline{\tau} - \tau_i$ refers to the difference between the average tax rate in the respondent's county and the tax rate paid by the respondent. In both panels the x-axis corresponds to the potential feedback that could have been shown to the subjects (i.e., the actual difference in tax rates). In panel (a) the y-axis corresponds to the *prior* beliefs about that difference (i.e., the respondent's perceptions before the feedback could have been shown) while in panel (b) the y-axis is the corresponding *posterior* belief (i.e., after the information provision experiment). The results from panel (b) are broken down by treatment group: the red diamonds (labeled "Feedback") correspond to respondents who were shown the feedback while the blue circles (labeled "No Feedback") correspond to those not shown the feedback.

Figure 6: Protest Rates and Attitudes, By Political Party

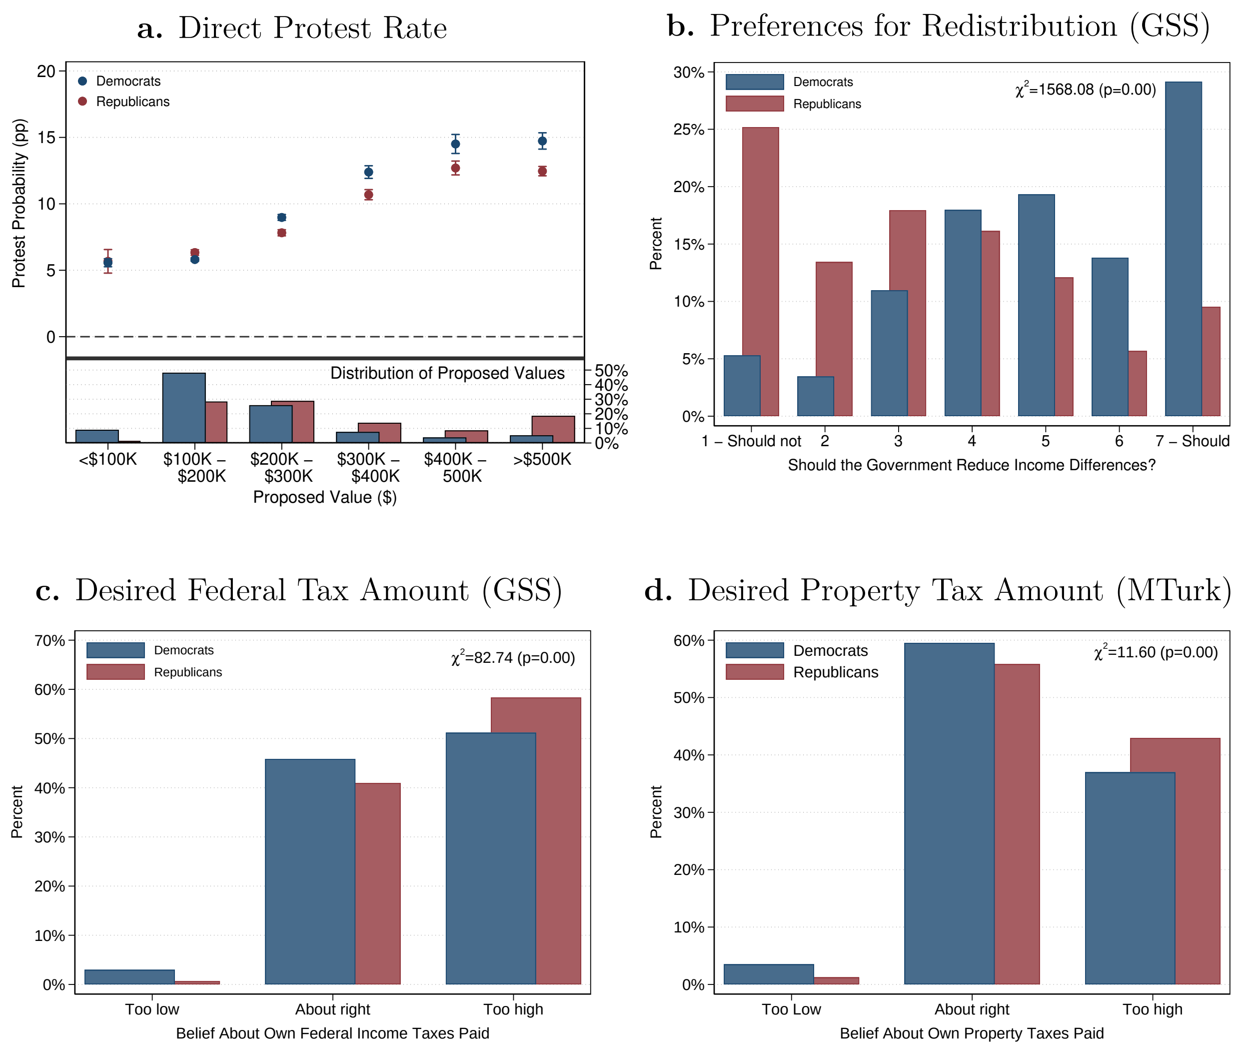

Notes: Panel (a): The direct protest rates for each proposed value bracket reported on the top of this panel and the share of households in each proposed value bracket reported on the bottom of this panel are based on the main sample of 423,607 single-family homes in 2020. The protest rates on the top of this panel represent point estimates with 95% confidence intervals in brackets, based on robust standard errors. Panel (b): This panel uses responses to the following question, broken down by self reported political party, from the General Social Survey (GSS): "Some people think that the government in Washington ought to reduce the income differences between the rich and the poor, perhaps by raising the taxes of wealthy families or by giving income assistance to the poor. Others think that the government should not concern itself with reducing this income difference between the rich and the poor. Here is a card with a scale from 1 to 7. Think of a score of 1 as meaning (...) What score between 1 and 7 comes closest to the way you feel?" Panel (c): this panel uses responses to the following question, broken down by self reported political party, from the GSS: "Do you consider the amount of federal income tax which you have to pay as too high, about right, or too low?" Panel (d): This panel uses responses to the following question, broken down by self reported political party, using data from the MTurk survey: "Do you consider the amount of property taxes you pay to be too low, about right, or too high?"

Figure 7: Effects of the Homestead Cap on Homestead Households: Heterogeneity by Political Party

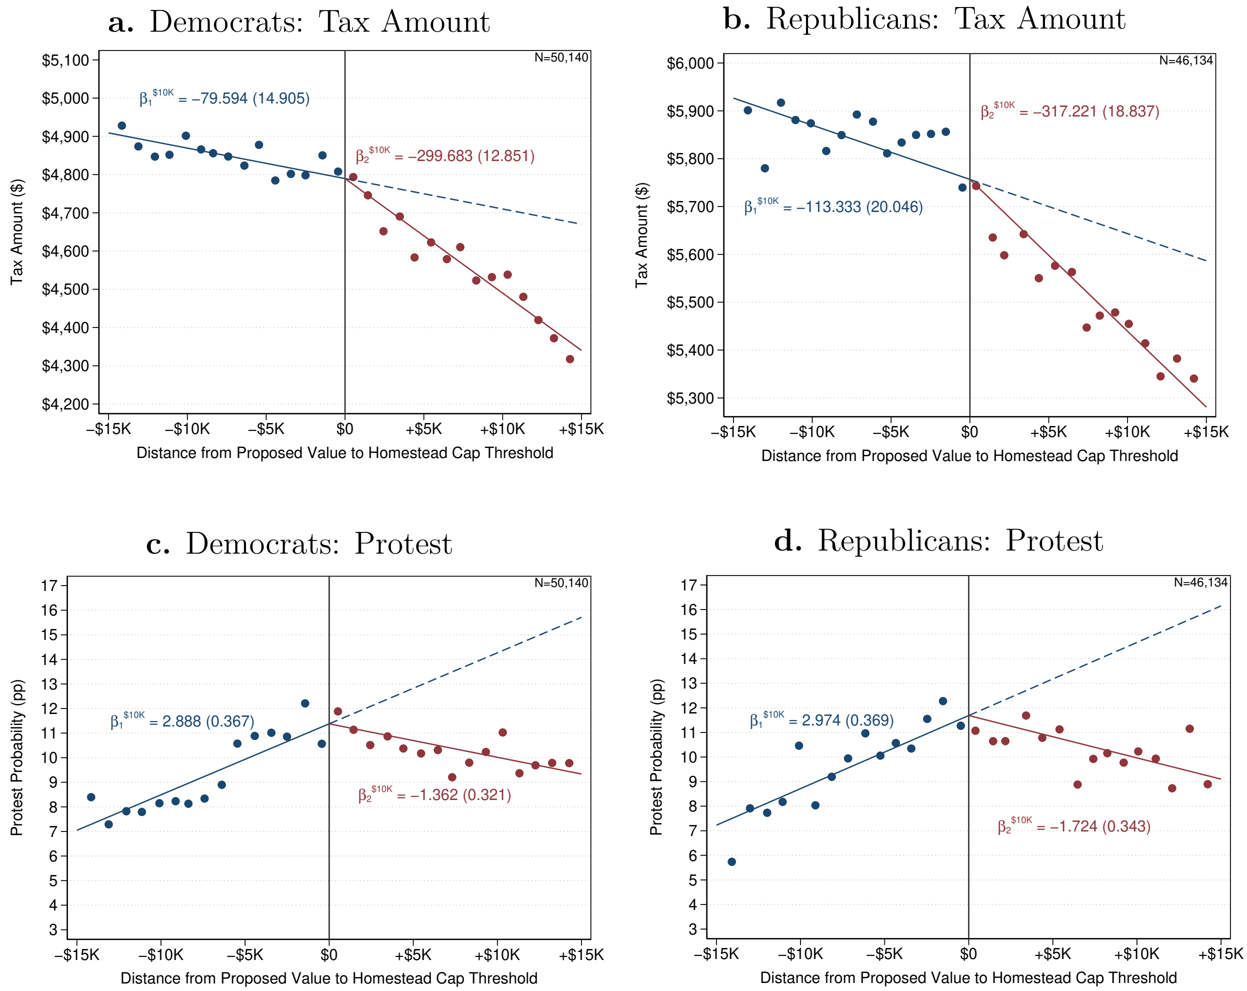

Notes: This figure features binned scatterplots of the relationship between a given outcome (indicated on the y-axis of each panel) and the distance between the 2020 proposed value and the 2020 homestead cap threshold (defined as 110% of the appraised value in the previous year). All regressions control for the proposed value, a dummy for whether the household protested in the previous year, and a set of school district dummies. The sample consists of single-family homes in 2020 for which the proposed value is within $15,000 of the potential homestead threshold. The lines correspond to linear regressions, with normalized slopes reported next to them along with robust standard errors (in parentheses) and the number of households (in brackets in the top right corner). The panels on the left half ((a) and (c)) correspond to households with 2020 homestead status who belong to likely Democrats, while the panels on the right half ((b) and (d)) correspond to households with 2020 homestead status who belong to likely Republicans. The dependent variables are: *Tax Amount* is the estimated tax amount based on 2020 proposed values, and P_{2020} is an indicator variable that takes the value 100 if the household protested directly in 2020 and 0 otherwise.

Table 1: Effects of the Two Types of Letters on the Probability of Protesting

	(1) P_{2020}	(2) P_{2020}^{won}	(3) P_{2020}	(4) P_{2020}	(5) SO_{2020}	(6) P_{2020}	(7) P_{2020}
Basic Aid Letter$^{(i)}$	1.792***	1.213***	1.449***	1.935***	0.795	1.943***	1.509***
	(0.249)	(0.222)	(0.347)	(0.339)	(0.719)	(0.317)	(0.394)
Extra Aid Letter$^{(ii)}$	3.509***	2.621***	3.108***	3.745***	15.287***	3.027***	3.994***
	(0.258)	(0.231)	(0.364)	(0.350)	(0.979)	(0.326)	(0.412)
P-value (i)=(ii)	<0.001	<0.001	<0.001	<0.001	<0.001	0.002	<0.001
Subsample			I	II		Dem.	Rep.
Mean Outcome (No Letter)	8.67	6.76	6.03	10.33	3.37	7.49	10.14
Std. Dev. Outcome (No Letter)	28.14	25.10	23.80	30.43	18.05	26.32	30.19
Observations	78,462	78,462	30,356	48,106	5,026	43,208	35,254

Notes: Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. Each column presents results from a different regression with two main independent variables: *Basic Aid Letter* is an indicator variable that takes the value 1 if the subject was randomly chosen to receive a basic aid letter and *Extra Aid Letter* is an indicator variable that takes the value 1 if the subject was randomly chosen to receive an extra aid letter. The omitted category is comprised by subjects who were randomly chosen not to receive a letter. The regressions in this table include the following controls: the proposed value in levels and its annual growth, dummies for multiple owners, school and special districts, number of years since the last protest, a dummy for homestead status, and for each year since 2015, a dummy indicating if the household protested in that year and the outcome of the protest (if any) as a %-reduction in the market value. The dependent variables are defined as follows: P_{2020} is an indicator variable that takes the value 100 if the owner filed a direct protest in 2020 and 0 otherwise; P_{2020}^{won} indicates with 100 if a direct protest resulted in a reduction in the assessed value, SO_{2020} is defined for the subsample that protested directly online and provided an opinion of value, and it takes the value 100 if the subject provided an opinion of value within half a percentage point of the value we selected for their extra aid message. Column (3) corresponds to the sample who were not mailed an official notification from the DCAD. Column (4) corresponds to the sample who were mailed such notification. Columns (6) and (7) split the sample between likely Democrat or Republican.

Table 2: Effects of the Information Revealing the Average Tax Rate in the County on the Probability of Protesting

	Mturk	Field Experiment							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	$\overline{\tau}_i^{post}$	P_{2020}	P_{2020}	I_{2020}	U_{2020}	P_{2020}	P_{2020}	P_{2020}	P_{2020}
Information Shock ($\overline{\tau}$)	0.393***	-0.950*	-15.392***	-13.220***	-0.468*	-11.050	-14.878**	-18.317**	-10.922
	(0.071)	(0.509)	(5.591)	(4.575)	(0.243)	(20.313)	(6.556)	(8.204)	(7.837)
Field Survey			✓	✓	✓	✓	✓	✓	✓
Subsample						I	II	Dem.	Rep.
Mean Outcome	1.24	11.29	50.26	81.90	7.12	50.30	50.25	50.47	50.10
Std. Dev. Outcome	0.77	31.65	50.01	38.52	2.15	50.04	50.02	50.03	50.02
Observations	2,065	50,394	1,888	1,867	1,888	672	1,216	860	1,028

Notes: Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. All columns present results from equation 7 in Section 5.3. The variable *Information Shock* ($\overline{\tau}$) corresponds to the information shock term ($D_i \cdot (\overline{\tau} - \tau_i)$). Column (1) reports results from the subjects in the Mturk survey, Column (2) from the subjects in the field experiment who received a letter, and Columns (3) through (9) from the subjects who received a letter in the field experiment and answered the Field survey. The dependent variables are defined as follows: $\overline{\tau}_i{}^{post}$ is the posterior belief on the average tax rate in the county; P_{2020} is an indicator variable that takes the value 100 if the owner protested directly in 2020 and 0 otherwise; I_{2020} is an indicator variable that takes the value 100 if the household reported to be either likely or very likely to protest in 2020 and zero otherwise; U_{2020} corresponds to a question about whether the taxes of the respondent are unfair relative to the taxes of everyone else, in a 1-10 scale. Columns (6) and (7) split the sample used in column (3) in two groups: i) subjects who protested in 2020 but did not protest in the recent past (column (6)) and ii) subjects who protested at least once during 2015 through 2019 or never protested during 2015-2020 (column (7)). Columns (8) and (9) split the sample in column (3) between likely Democrat or Republican.