

NBER WORKING PAPER SERIES

MY TAXES ARE TOO DARN HIGH:
WHY DO HOUSEHOLDS PROTEST THEIR TAXES?

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 May 2022

We are thankful for excellent comments from Youssef Benzarti, Emiliano Huet-Vaughn, Justin Holz, John List, Erzo Luttmer, Michael Norton, Stefanie Stantcheva, Matt Weinzierl, and other colleagues and seminar participants at the NBER-Public Economics, Stanford University, University of Chicago, Yale University, Universidad de San Andres, the American Economic Association, NTA, ZEW, Journees LAGV, LMU, RIDGE, and AEI. This project was reviewed and approved in advance by the Institutional Review Board at The University of Texas at Dallas. The experiments were registered in the AEA RCT Registry (#0005992). The original pre-registration was posted on May 24, 2020. However, after receiving interview requests from the media, we removed the pre-registration (we were concerned that the media would divulge the hypotheses listed in the pre-registration and thus contaminate the study). We re-posted the pre-registration on June 16, 2020, after the deadline to file a protest had passed (and before the data to conduct the analysis became available). After the study is accepted for publication, we will make all the code and data publicly available. Adrian Cadena Medina, Luisa Cefala, Dongwook Chun, Karl Dill, Santiago De Martini and Xinmei Yang provided superb research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Brad C. Nathan, Ricardo Perez-Truglia, and Alejandro Zentner. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

My Taxes are Too Darn High: Why Do Households Protest their Taxes?
Brad C. Nathan, Ricardo Perez-Truglia, and Alejandro Zentner
NBER Working Paper No. 27816
September 2020, Revised May 2022
JEL No. C93,H2,H26,Z13

ABSTRACT

In the United States and many other countries, taxpayers can file a protest to legally reduce their property taxes. While tax protests can provide a unique opportunity to study the (un)willingness to pay taxes, they have received little attention from researchers. To fill that gap, we study what motivates households to protest their property taxes. Using a field experiment, we show that filing frictions and fairness considerations play significant roles. In comparison, partisan identity plays a minor role. We calculate the magnitude of filing frictions and willingness to pay for fairness using a money metric. To do so, we combine our experimental estimates with quasi-experimental evidence on the role of expected tax savings. Last, we discuss how low-cost interventions targeted at disadvantaged groups can mitigate existing economic and racial disparities in the system of tax appeals.

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>
A online appendix is available at <http://www.nber.org/data-appendix/w27816>

1 Introduction

Property taxes are an important source of revenues for governments in the United States and around the world. For example, property taxes are the second largest source of tax revenues in the United States, generating an estimated \$547 billion in 2018 (Urban Institute, 2021).¹ In all U.S. states, as well as in many other countries, households can file protests to *legally* reduce the amount they have to pay in property taxes.² These protests are consequential: while there is no guarantee, tax protests can substantially reduce the tax amount due. Tax protests provide a unique opportunity to study the (un)willingness to pay taxes, yet this context has received little attention from researchers. To fill that gap, we take the first steps by disentangling households’ motivations for filing a property tax protest.

This paper introduces a new high stakes naturally occurring setting that can be used to conduct field experiments. We conducted our field experiment in Dallas County, Texas. Tax protests in Dallas County work in an almost identical way in all 241 counties in Texas,³ and, while some differences exist, there are similarities in how they work in other U.S. states and countries.⁴ We focus on a single county purely for the logistical advantage from implementing a field experiment in a single location. Dallas County is the second largest county in Texas, with an estimated population of about 2.6 million in 2020 – indeed, Dallas County has a larger population than 15 of the 50 U.S. states (U.S. Census Bureau, 2021). The county is diverse along many dimensions including, but not limited to, ethnicity and political party representation. Texas does not have a state income tax and therefore property taxes are a key source of revenue for the provision of government services. The average household in Dallas County was expected to pay around \$5,916 in property taxes in 2020, corresponding to an average tax rate of 2.01% of home value.⁵

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

¹ For reference, the federal income tax generated \$1.7 trillion and the corporate income tax \$230 billion in 2019 (Stein and Ingraham, 2019).

² See for example Dobay et al. (2019), who overview how property tax appeals work across states and provinces in the United States, Canada, and Australia and at the country level in Hong Kong, Ireland, New Zealand, South Africa, and the United Kingdom. World Bank (2019) overviews how appeals work across Latin American and Caribbean Countries.

³ For a history of property tax recapture in Texas, see for example Villanueva (2018).

⁴ For example, property owners must pay a filing fee in Alaska. Due to institutional factors, tax protests are less common in some contexts (e.g., in California appraised values are updated only when properties are sold and thus there is much less of an incentive for households to file a protest).

⁵ These statistics are based on administrative data focusing on single-family households. Throughout this study, we use the term “tax rate” to refer to a household’s effective tax rate (computed as the household’s total property tax amount billed 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; we refer to this amount as the *proposed value*. Property taxes are calculated based on the proposed value, and the DCAD notifies the household of the amount. The homeowner has the option to file a protest, for example, arguing that the proposed value (and thus the corresponding tax due) is too high.⁶ Owners can protest directly on their own (which is the main focus of this paper) either online or by mail, or they can hire an agent to protest on their behalf. The DCAD then responds to the homeowner’s appeal. We refer to a protest as being successful if the DCAD lowers the effective assessed value. In 2020, 8.40% of households in Dallas County filed a protest on their own and an additional 8.42% of households protested with the help of an agent, resulting in a total protest rate of 16.82%. We estimate that 69.7% of protests were successful in 2020, resulting in average tax savings of \$485 (8.2% of the average property tax bill) in the first year alone.⁷

The average protest rate masks large heterogeneity. For example, wealthier homeowners have a much higher probability of protesting: the average protest rate is 42.0% for homes worth over \$500,000, but only 8.9% for homes worth less than \$100,000. In addition, among homes of similar value, there are significant racial and demographic differences in the probability of protesting. For instance, Hispanic households are significantly less likely to file a protest than their White counterparts; and older homeowners are much less likely to protest online than their younger counterparts and more likely to protest by mail. The reasons for these differences are unclear. Some individuals have been quite critical, arguing that the “property tax system is rigged against (...) little people” (Lieber, 2020). This interpretation assumes that some groups are left behind because, for example, they cannot figure out how to protest, they are less technologically savvy, or they cannot afford a tax agent (Doerner and Ihlanfeldt, 2015). However, these differences may be due to other factors. For example, wealthier households may protest more because they feel that they are paying more than their fair share. If we ever want to provide useful insights for the design of an equitable appeal process, we first need to address the most basic question: why do households choose to protest their taxes?

We start by measuring the extent to which expected tax savings affect the decision to file a tax protest. While the direction of households’ responses to this basic economic incentive is hardly surprising, quantifying this effect plays a key role in that it serves as a money metric: i.e., it allows us to translate other estimates, such as filing costs and fairness concerns, to dollar amounts. To study the role of expected tax savings, we exploit quasi-experimental

⁶ We describe the protest process in more detail in Section 2.

⁷ These estimates correspond to owners who protested directly in 2020. These and other estimates in the paper are based on the administrative data as of November 2020 and thus do not include a few protests that may have been resolved after November 2020.

variation introduced by a feature 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.⁸ The homestead cap generates a sharp kink in the marginal benefits from protesting. When the proposed value is below the threshold, a marginal reduction in the proposed value reduces the amount due in property taxes. However, when the proposed value exceeds the threshold, a marginal reduction in the proposed value has no effect on the tax amount. Exploiting this exogenous variation, we find that households are indeed very responsive to their expected benefits from filing a tax protest. Specifically, a \$100 increase in the marginal benefits from protesting causes an increase of 2.14 percentage points (pp) in the probability of protesting.

We designed a field experiment to explore the role of filing frictions and fairness considerations in the decision to protest taxes. We conducted the field experiment with a subject pool of 78,462 households. The first treatment arm studies filing frictions. In Dallas County, any household can protest property taxes for free, meaning that protesting incurs no pecuniary costs. However, households may face filing frictions. These frictions include the hassle costs from filing taxes, such as the opportunity cost of time and the unpleasant nature of doing paperwork (Goolsbee, 2006; Benzarti, 2020; Sunstein, 2021; Benzarti, 2021). Additionally, our definition of filing frictions includes information frictions. For example, some households may not protest because they do not know how to protest, or because they believe filing a protest is a lot harder than it really is. Of the 78,462 households in the subject pool, 50,394 were randomly assigned to receive a letter aimed at reducing filing frictions. We randomized each of these households into one of two letter types, with increasing degrees of treatment intensity. The *basic aid letter* included a step-by-step guide for filing a protest by mail or online. The *extra aid letter* included the same information as the basic aid letter as well as additional instructions on how to complete one of the most challenging aspects of the process: preparing an argument to support the protest.⁹

The evidence from the first treatment arm indicates that filing frictions are of first order importance in the decision to protest taxes. We find that receiving a letter had a large positive effect on the probability of filing a protest, and that the letter that offered more aid induced a larger effect. These effects are not only highly statistically significant, but economically significant too. For example, the extra aid letter increased the protest rate by

⁸ Homestead is a legal status that can be granted to a property that is the owner's (owners') primary residence. This status comes with several advantages related to property taxes (e.g., exemptions and caps) and other benefits (e.g., exemption from forced sale for collection of debt).

⁹ In the extra aid message, we included an argument tailored to each recipient that could be copied into the protest form. These letters presented information about another property near the recipient's own property that was comparable in all observable characteristics and was recently sold for a lower price than the market value proposed by the DCAD.

4.98 pp.¹⁰ For reference, this effect is equivalent to 57.4% of the baseline protest rate.¹¹

We use multiple strategies to show that the letters worked through reducing hassle costs and information frictions, rather than merely making the protest option more salient. Moreover, to quantify the magnitude of the filing frictions, we combine the experimental and quasi-experimental estimates. Our back-of-the-envelope calculations indicate that filing frictions are on the order of \$232. Indeed, this estimate constitutes a lower bound, as our intervention is probably far from reducing the filing frictions in full. This result suggests that filing frictions are a key reason why some households do not file a tax protest. As a result, filing frictions should be taken into account in the design of a more equitable system of tax appeals.

In the second treatment arm, we explore the role of fairness considerations (see e.g., Andreoni, 1995; Fehr and Schmidt, 1999; Fehr and Schurtenberger, 2018). More specifically, we explore the conditional cooperation channel: i.e., whether households are more willing to pay taxes if they think that other households are paying their fair share too. If households feel they overpay relative to other households they may be more likely to file a protest. To motivate the design, we start with some non-experimental data. First, we show that households who say that their taxes are more unfair are significantly more likely to file a tax appeal.¹² Second, we provide survey evidence that a strong majority of homeowners believe that every household should be subject to the same tax rate.

Our experimental design leverages misperceptions about the average tax rate: homeowners who pay below-average tax rates tend to underestimate the average tax rate, and homeowners who pay above-average tax rates tend to overestimate the average tax rate.¹³ We designed an information-provision experiment that leverages these systematic biases to create exogenous variation in the perceived average tax rate. In the letter sent to each household, we randomize whether, in addition to the information on the household's own tax rate, information on the average tax rate paid in the county is also included. We show that the information shocks induced by the experiment have significant effects on households' beliefs and behaviors, and in the directions predicted by the conditional cooperation channel. More specifically, when a household is shown that the average tax rate is higher than originally thought, that causes the household to: (i) perceive its own taxes as more fair (according to their answers observed in the survey data); (ii) reduce the probability of filing a tax protest

¹⁰ This is our preferred treatment-on-the-treated estimate, which accounts for some letters failing to be delivered or not being read by the subjects. In comparison, the raw intention-to-treat effect is 3.51 pp.

¹¹ An 8.67% of households in the control group, which did not receive any letter, filed a protest.

¹² This finding is based on survey responses from 1,888 subjects from the field experiment. More precisely, households were asked to indicate whether their own tax rate is fair or unfair, relative to the county average.

¹³ This finding is based on a supplemental online experiment with 2,065 respondents as discussed in Section 5.4.

(according to their behavior observed in the administrative data).

We use multiple strategies to show that the information about the average tax rate operated through the fairness channel, rather than potential alternative mechanisms.¹⁴ Moreover, we show that the fairness channel is not only statistically significant, but also significant in magnitude. Our estimates suggest that a 0.1 pp increase in the perceived average tax rate decreases the protest probability by 3.9 pp. Combining our quasi-experimental estimates on expected tax savings and our experimental estimates on conditional cooperation, we calculate households' willingness to pay for fairness. We estimate that a 0.1 pp increase in the perceived average tax rate increases the mean household's willingness to pay taxes by about \$182.24.

As a benchmark for the above findings, we explore an alternative mechanism that is widely discussed in both academic and non-academic audiences: partisan identity. According to survey data and laboratory experiments, there are some significant differences in attitudes towards taxation between Republicans and Democrats (Huet-Vaughn et al., 2019; Stantcheva, 2020). We categorize households as Democrat or Republican by matching the individuals from the property tax rolls to other sources of data, such as the Texas voter files. We show that, relative to hassle costs and fairness, partisanship plays a minor role in explaining the decision to protest taxes. Some differences along party lines are statistically significant, but they are mostly small in magnitude.

The use of property tax protests is a novel and promising high-stakes setting to conduct natural field experiments for several reasons. First, in addition to the questions studied in this paper, property tax protests provide an alternative setting to study both new behavioral questions and questions previously analyzed using other settings, such as those studied in collaboration with the tax authority or charitable organizations, or using public data (e.g., voting).¹⁵ Second, this setting allows interventions based entirely on publicly available data which allows replicability. Last, the use of this setting does not require a partner organization nor data agreement, which allows studying questions, and avoiding potential publication restrictions of findings, that do not align with the interests of partner organizations.

Our study contributes to the understanding of tax appeals. Even though tax protests are common in the United States and worldwide (Dobay et al., 2019; World Bank, 2019), they have received little to no attention in the economics literature. In comparison, other forms

¹⁴ For example, Section 5.6 provides multiple pieces of evidence against a potential confounding factor: that households use the information on average tax rates to infer something about the probability that they will succeed in their protests.

¹⁵ For instance, the setting could be used to examine how taxpayers react to nudges as studied in the tax evasion literature, aspects related to altruism or warm-glow as studied in the setting of charitable giving, or aspects related to behavior induced by incentives “to tell others” as studied in the setting of voting in elections.

of tax compliance, such as tax evasion and tax delinquency, have been studied extensively (Luttmer and Singhal, 2014; Slemrod, 2019). One notable exception is Jones (2019), who uses data on the decision to protest taxes to provide a test of loss aversion.¹⁶ Another notable exception is Avenancio-León and Howard (2019), who show that property taxes are disproportionately higher for racial and ethnic minorities and that some of those differences may be related to the system of tax appeals.¹⁷ We contribute to this literature by using experimental and quasi-experimental variation to measure the motives behind the decision to file a tax protest. These motives are key inputs for a discussion on the design of an equitable tax appeal system. Indeed, our findings have some direct policy implications. For example, we show that low-cost interventions targeted at disadvantaged groups can go a long way in mitigating the unequal access to the protest system.¹⁸

Our study is also related and contributes to other strands of literature. Our findings on fairness considerations contributes to the literature on tax morale (Luttmer and Singhal, 2014; Slemrod, 2019). Our contribution to this literature is twofold. First, we are able to provide unique causal evidence that fairness considerations matter for tax compliance. Indeed, in a follow-up study (Giaccobasso et al., 2022), we leverage the context of tax protests to provide further causal evidence on the importance of tax morale. Second, we measure the willingness to pay for fairness with a money metric. Our findings are also related to a recent literature on tax filing costs (Benzarti, 2020, 2021). We contribute to this literature by providing experimental evidence as well as measuring the filing frictions with a money metric. Our findings are also related to recent literature on partisan polarization (Stantcheva, 2020; Huet-Vaughn et al., 2019).¹⁹ We contribute to this literature by providing experimental evidence from a natural context and with high-stakes behavior. We also show that, in contrast to widespread perceptions on polarization (Ahler, 2014), Democrats and Republicans behave quite similarly in our setting.

The rest of the paper proceeds as follows. Section 2 describes the institutional context. Section 3 studies the role of expected tax savings. Section 4 studies the role of filing frictions. Section 5 discusses fairness considerations. Section 6 analyzes the heterogeneity by political

¹⁶ 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 upward revisions than for the downward revisions.

¹⁷ Some other studies look at property taxes more generally, without focusing on protests. For example, Cabral and Hoxby (2012) provide evidence on how the salience of property taxes can affect equilibrium tax rates. Tax appeals have also received the attention of legal scholars (Hayashi, 2014).

¹⁸ This view is consistent with related evidence from other contexts such as social benefits (Alatas et al., 2012, 2016; Finkelstein and Notowidigdo, 2019) and income tax filing (Bhargava and Manoli, 2015; Benzarti, 2020).

¹⁹ This literature relies heavily on evidence from surveys and laboratory experiments. One notable exception is Cullen et al. (2020): using aggregate statistics from the tax records, they show that taxpayers evade federal taxes less often when the president belongs to their own political party.

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 the 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.²⁰ 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 protest 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

²⁰ In Dallas County, the four county-level taxing jurisdictions are the county (including a community college, a hospital, and the county itself), 31 cities, 16 ISDs and 12 Special Districts (Source: <https://comptroller.texas.gov/taxes/property-tax/county-directory/dallas.php>). Appendix A.1.1 provides additional details about the property taxes in Dallas County and the data sources.

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 filing frictions 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. Moreover, homeowners need to file a form in order to remove an agent, which creates a stickiness and implies that agents often protest every year on behalf of owners.²¹ 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 a fraction of households, by mailing a “Notice of Appraised Value”.²² 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.²³ In 2020, the DCAD presented the proposed values on May 15th; as a result, the deadline to protest was June 15th.

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

²¹ 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.

²² Notifications are mailed to households meeting certain criteria such as increased appraised value, ownership change, loss of homestead exemption, rendered property, or new property.

²³ 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).

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. Indeed, the challenge to estimate home values has recently brought one of these companies financial trouble (Parker and Friedman, 2021). This ambiguity in market values leaves room for the owners to complain about the DCAD’s value assessments.²⁴ 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: using an independent estimate of household market value (Redfin), households still file protests when their properties have been under-assessed by the government.²⁵ This view is also consistent with responses to open-ended questions from households in the sample. Our Field Survey (introduced in Section 4.2 below) included an open-ended question on why the household will, or will not, protest in 2020. A majority (55.9%) of households mentioned that their taxes were too high.²⁶ For example, one household explained that it would protest “because my taxes are way too high.”

Filing a protest directly is simple. Households can protest using either 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, about 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

²⁴ 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.

²⁵ In Appendix A.1.4 we show that protests increase with the difference between the Redfin estimate and the proposed value, but that homeowners whose proposed values are under the Redfin estimate still file protests.

²⁶ These statistics are computed from hand-coding the answers to the open-ended question included in the survey, as a fraction of the non-missing responses. There were a total of 1,116 respondents mentioning that their taxes were too high.

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.²⁷ The formal hearing entails no risk: if the DCAD schedules a hearing and the household does not attend, the protest is simply dismissed with no penalty. Of the 2020 direct protests that contain information on the form in which it was resolved, we find that 43.7% were resolved informally, 35.2% were resolved after a formal hearing, and 21.1% were either withdrawn or dismissed. 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.²⁸ Our survey included an open-ended question on the reasons why households choose whether to protest or not. Of all respondents, a small minority (3.9%) mentioned that their decision on whether or not to protest was related to the pandemic. 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.²⁹ 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.³⁰ 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 selected to participate in the field experiment.³¹

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.³² If we include protests through agents too, the protest rate is 16.83%. This rate of protests has been quite stable in recent

²⁷ Homeowners can rescheduled the hearing one time without demonstrating any cause, including those that miss a scheduled hearing if the request is made within 4 days of the missed scheduled hearing. 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.

²⁸ If all else fails, the homeowner has the option to contest the decision in court.

²⁹ For more details on the definition of this sample, see Appendix A.1.3.

³⁰ 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.

³¹ In Appendix A.1.3 we provide descriptive statistics for each of the subsamples used in the study, and show that they are roughly similar in a number of key respects.

³² In all the analyses presented in this paper, we include protests that were marked as received by the DCAD through July 15th, 2020. For more details, see Appendix A.1.2.

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 over 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 2015–2019.³³

The above probabilities are based on averages, and as such they mask substantial heterogeneity across households. These sources of heterogeneity are important in that they can reflect unequal access to the system of tax appeals. For example, wealthier homeowners have a much higher probability of protesting. And even within households of similar wealth, certain groups of the population are systematically more likely to protest than other groups. For example, conditional on owing the same amount in taxes, Hispanic households are 3.61 pp less likely to protest directly than their White counterparts. Likewise, households with members 40 years old or younger are 5.22 pp more likely to protest directly than comparable households with members aged 65 or above. Moreover, older homeowners are far less likely to protest online than their younger counterparts and more likely to protest by mail, suggesting a digital divide.³⁴

3 Expected Tax Savings

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 order to estimate the expected tax savings from protesting. 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, we now introduce a simple model of the decision to protest. Let A be the proposed value of the

³³ For more details, see Appendix A.1.4.

³⁴ See Appendix A.1.4 for details on these auxiliary statistics.

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:

$$T_{nocap} = \tau \cdot A \quad (1)$$

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:

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

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 homestead cap threshold be \bar{A} . Taking this threshold into consideration, the tax burden can be computed as follows:

$$T_{cap} = \tau \cdot \min\{A, \bar{A}\} \quad (3)$$

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:

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

This equation can be re-arranged as follows:

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

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 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 were approved for homestead status in 2020, and thus their homestead caps may be binding. The remaining 26% of 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 also is reported in blue. For ease of exposition, we normalize all 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. The results from Figure 1 includes a set of basic characteristics of the household as control variables: 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. The two panels on the right of Figure 1 (1.b and 1.d) correspond to properties without homestead status, for which the homestead cap threshold should be irrelevant and thus they can 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.³⁵ The blue slope in the left half of the figure corresponds to properties right below the homestead cap threshold, and the red slope in the right half of the figure corresponds to the properties right above the threshold.³⁶ Note that the slope on each side represents just an association and thus only the difference between the two slopes can be interpreted as a causal effect. As expected, there is

³⁵ 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.

³⁶ As a robustness check, Appendix A.2.3 shows that there is no bunching at the homestead cap.

a sharp kink at the threshold: upon reaching the homestead cap threshold, households pay a lower tax amount than they would have without the homestead cap. This kink is large in magnitude and statistically significant. We therefore 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.

The following thought experiment illustrates 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, 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 (= -306 - (-97)), corresponds to the tax amount that is capped.³⁷ What happens when a successful protest decreases the proposed value by \$10,000? In the presence of the homestead cap, the \$10,000 reduction in household value would not affect the tax amount of the household. That is, in the presence of the homestead cap, the household would not benefit from the \$209 reduction in taxes because the household already benefited from that reduction due to the homestead cap.

In sum, for a household at +\$10K on the x-axis of Figure 1.a, the marginal benefit from protesting decreases by \$209. If households care about the expected tax savings from protesting, we expect the \$209 reduction in the marginal benefit from protesting to reduce the probability of protesting. To explore this expectation, 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). As expected, we find a sharp kink in Figure 1.c at exactly the homestead cap threshold. Again, this kink is large in magnitude and 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 an average \$209 reduction 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 pp. 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. In other words, the decision to protest is responsive to the expected tax savings from protesting. In Sections 4 and 5, we will use this estimate to calculate money metrics for both the filing frictions and the willingness to pay for fairness.

³⁷ 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

The previous analysis corresponds to properties that have a homestead status exemption (and thus are subject to the homestead cap). Next, we reproduce the analysis using properties that do not have a homestead status exemption (and thus are not subject to the homestead cap). This analysis provides a sharp falsification test because we should not observe any kinks in this latter group. A kink would suggest that the effects are not due to the homestead cap but instead due to some other confounding factors. The two panels in the right half of Figure 1 (i.e., 1.b and 1.d) correspond to the properties without homestead status for which the hypothetical homestead cap threshold is defined as 110% of the assessed value in the previous year (2019). As expected, on the right-side panels of Figure 1, we find no kinks at the homestead cap threshold; thus, 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 on the right-side panels of Figure 1, meaning that we can rule out any large kinks shown in the left-side panel of Figure 1, as well as small kinks.

For a second falsification test, we follow the logic of an event-study analysis. We reproduce the analysis from Figure 1.c for the properties with homestead status, but the dependent variable is whether the household protested in 2019 instead of 2020. Whether the 2020 proposed value ends up above or below the 2020 homestead cap should not affect whether a household protested in 2019. In fact, a significant difference would suggest a confounding factor affecting the results from Figure 1.c. Moreover, we can extend this logic and reproduce the analysis for one year prior (2019) and for each year for which we have data (2015–2018). Figure 2.a presents the results from this event-study analysis. 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 different year. For example, the 2015 coefficient corresponds to a regression where the dependent variable takes the value 100 if the household protested in 2015 and 0 otherwise.³⁸ As expected, the coefficients are consistently close to 0 for each falsification year (2015–2019) and always highly statistically different from the coefficient for 2020.³⁹

In Appendix A.2, we present several additional robustness checks. We show that, in addition to affecting direct protests, the homestead cap affects agent-assisted protests in

³⁸ Note that the set of control variables will not be identical. For example, 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, we control for whether the household protested in 2018.

³⁹ 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 expect a few to be statistically significant just by chance.

the same direction, but the 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. The results are qualitatively and quantitatively similar if we use alternative bands. We also show that the homestead cap is consequential not only for the number of protests but also for the subsequent market values and tax amounts. Last, we show that there is no bunching at the homestead cap.

4 Filing Frictions

4.1 Conceptual Framework

There are no fees for filing a property tax protest in Texas. However, we hypothesize that households may face non-pecuniary costs, which we call filing frictions. These frictions include the traditional hassle costs from filing taxes such as the opportunity cost of time and the unpleasant nature of doing paperwork (Goolsbee, 2006; Benzarti, 2020; Sunstein, 2021; Benzarti, 2021). Additionally, our definition of filing frictions includes information frictions. That is, some households may not protest because they do not know how to protest or because they think that filing a protest is a lot harder than it really is.

While filing a protest directly is easy in theory, for some households it may not be easy in practice. Some households may not even know where to start, or they may think the process is 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 DCAD website.⁴⁰ However, this long document includes broad instructions and is difficult to locate on the DCAD website. Other unofficial online sources, such as blog posts, 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 filing frictions may be particularly large for a specific step in the protest procedure: providing an opinion on the value of the home and an argument supporting that opinion. This information can be obtained in different ways, according to anecdotal evidence, and usually involves a comparable property that recently sold for less than the proposed value of the contested property. The recent sale price serves as the opinion of value, and information about the recent transaction can be used as the argument. Finding a proper comparison property entails several steps. First, the homeowner could use a free online real estate

⁴⁰This document can be found at http://www.dallascad.org/Forms/Protest_Process.pdf.

service, such as Zillow.com or Redfin.com, to search for and identify comparable properties that sold recently for less than the homeowner’s proposed property value. This seemingly straightforward process could be daunting for people with limited Internet access and 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 filing frictions. Subjects in our sample were randomly assigned to receive a letter or no letter. Figures 3 and 4 show the first and second pages of a sample letter, with the addition of red boxes highlighting the parts that were randomized.⁴¹ The letters were sent on behalf of researchers at The University of Texas at Dallas and included several measures to show that they came from a legitimate source.⁴² Letters were tailored to recipients: each salutation at the top of the first page included the recipient’s name, and their names and addresses were printed at the bottom of the second page (which appeared through the envelope window).⁴³

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* provided useful tips to help the recipient file a protest, all of which were printed on the *first* page of the letter (see Figure 3 for a sample). A key part of the first page is that it included instructions on how to file a protest using the project’s website, which we designed to be concise, easy to follow, and as explicit as possible. Appendix D shows screenshots of the entire website, including the step-by-step instructions on how to file a protest online or by mail. These walkthroughs included hyperlinks to relevant websites and screenshots of a sample protest using information from a fictitious household for added clarity.

The *extra aid letter* is identical to the basic aid letter, plus additional guidance on how to protest. Figure 4.a shows an example of the second page of the letter for a recipient assigned to the basic aid treatment, and Figure 4.b shows an example of the second page for a recipient assigned to the extra aid treatment, with the extra aid message outlined by a red box with

⁴¹ See Appendix C for a full-page sample of the letter without the red boxes added.

⁴² The envelope featured the logo of The University of Texas at Dallas, a well-known institution in Dallas County; the name of one of the professors from that university; and non-profit organization postage. See Appendix B for a sample. The letter also featured the official logo of The University of Texas at Dallas in the header, as well as a physical address and the URL of the study’s website. The website provided general information about the study (without discussing any hypotheses or what the study was about) and contact information for the researchers and the institutional review board.

⁴³ In cases where properties were jointly owned by multiple individuals (typically, husband and wife), we sent a single letter addressed to all listed owners.

dashed lines (we added this box for expositional purposes and did not include it in the actual letters sent to subjects). The extra aid message provided the additional information related to the opinion of value and supporting argument.

The first paragraph of the extra aid message offered some facts about the protest filing process, explaining that the simple process could be done without an agent and may not require a hearing (which could be intimidating to some subjects) for the DCAD to propose a settlement offer. Moreover, if a hearing were to be scheduled, there would be no risk for not attending. The message also provided an argument to be used in the protest. Specifically, we presented the most common approach, that is, arguing that based on a recent sale price of a comparable property, the proposed value for the property exceeds the market value.⁴⁴ To further simplify the use of this information, we presented this information as it would look on the actual protest form, with a check mark in the “Value is over market value” box, the comparable sale price in the “Opinion of value” field, and a sample handwritten note outlining an actual argument they can use in their protest. 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 use our proposed argument verbatim, but to clarify that our content was 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 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 Redfin.com. For each recipient, the algorithm searched similar properties (e.g., same number of bedrooms, bathrooms, square footage, location) that were sold in late 2019 or early 2020 for between 5% and 20% less than the proposed value of the recipient’s property.⁴⁵ We believe that an agent hired by the recipient to protest would follow a similar (or even the exact same) procedure.

Our letters were not designed to eliminate the filing frictions fully, as recipients would still need to spend time filing and monitoring their protests, which has an opportunity cost.

⁴⁴ We identified one comparable property for all households in the subject pool, but we included this information only for recipients who were randomly selected for the extra aid letter.

⁴⁵ In Appendix A.3, we provide details about this algorithm and some descriptive statistics.

Some also may find the associated paperwork to be considerably unpleasant (Benzarti, 2020, 2021). In that sense, our estimates provide a lower bound of the full magnitude of the filing frictions.

Some additional features of the letter are summarized here but discussed in detail in Section 5 where we discuss fairness concerns. In the middle of the first page, all letters included a table with information related to the proposed values and estimated taxes of the recipient’s property. At the bottom of the first page, all letters included a URL to an online survey (hereafter, we refer to this as the *Field Survey*). To verify that recipients were legitimate subjects and to link survey responses at the household level, we included a unique five-letter code for survey access. The main goal of the Field Survey is to include questions to be used as outcomes in the analysis discussed in Section 5. It also provides a proxy for the dates that recipients opened the letters (Perez-Truglia and Cruces, 2017; Bottan and Perez-Truglia, 2020a).

4.3 Subject Pool and Implementation Details

From the main sample of 423,607 residential single-family properties, we focused on a subgroup of 78,462 homes for our field experiment subject pool. To arrive at this subsample, we excluded households that had already filed a protest before receiving our letters according to the latest available data from DCAD because our letter could not affect their behavior.⁴⁶ The most important condition was to focus on households for whom our algorithm could find comparison properties for use in the extra aid message.⁴⁷ Though not identical, the subject pool is similar to the main sample in observable, pre-treatment characteristics.⁴⁸ 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 before the protest deadline to influence the recipient’s decision. We created the letters as soon as the administrative data, including 2020 proposed values, became available (on May 16th, 2020). To accelerate delivery, we used a mailing company in Dallas County (i.e., the same county as all recipients). The mailing company dropped off the letters at the local post office on May 20th and estimated that most would be delivered in the next couple of days. Consistent with this projection, we began to receive Field Survey responses and website visits on May 21st. Moreover, we confirmed that the post office scanned more than 90% of our letters by Friday,

⁴⁶ We initially selected a sample of 79,322 properties. However, due to a lag of a few days in the way DCAD reports data, we dropped 860 from the subject pool because we discovered that they had already protested by the time we mailed the letters.

⁴⁷ For more details, see Appendix A.1.2.

⁴⁸ See Appendix A.1.3.

May 22nd, 2020, indicating they reached the last mile before delivery. Based on data from previous years, most subjects file protests close to the deadline, which in 2020 was June 15th. Thus, we feel confident that there was enough time between receipt of the letter and the protest deadline for the letter to influence most recipients’ decisions to protest. Indeed, this view is consistent with the dates when subjects responded to the Field Survey and when they visited the website included in the letter.⁴⁹

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 the 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:

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

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. Last, 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 property tax protest, a dummy for homestead status, growth in the proposed value relative to the previous year and for each year from 2015 to the previous year, and dummies indicating whether the household protested in each 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

Table 1 presents the regression results. All regressions are based on the same specification given in equation (6), but they differ by 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

⁴⁹ Results reported in Appendix A.4.1.

significant (p-value<0.001). 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 coefficients show the effects of each type of letter on the probability of protesting in 2020, and they 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 the dependent variables are protest indicators for the years 2015 through 2019, instead of 2020. As our letters had not been sent yet, they could not possibly affect protests in prior years. As expected, the coefficients for the pre-treatment years are close to 0, 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., the protest was unsuccessful or the household did not protest directly). The coefficients remain economically and statistically significant. The ratio between the coefficients on the extra aid letters in columns (1) and (2) suggests that 75% ($= \frac{2.621}{3.509}$) of the marginal protests that were induced by our extra aid letter were successful. This success rate is comparable to the corresponding success rate of 78% ($= \frac{6.76}{8.67}$) observed in the control group.⁵⁰ In other words, the additional protests induced by our letters were, on average, roughly as successful as the protests in the control group.

In Appendix A.4.2, we present several additional robustness checks. We show that the increase in direct protests induced by our letters did not crowd out protests through agents as expected because, as explained before, households often sign long-term contracts with agents years or months before the proposed values are announced. We also 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: reduction in market value and reduction in estimated taxes.⁵¹

4.6 Causal Mechanisms of the Effects of the Aid Messages

In this section, we unpack evidence related to the causal mechanisms behind the effects of the aid messages.

Our favorite interpretation is that the letters increased the likelihood of protesting by

⁵⁰ 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 households that protested directly (8.67, from column (1)).

⁵¹ Results reported in Appendix A.4.2.

mitigating the underlying filing frictions. One interpretation could be that our letters reduced the filing frictions by acting as a reminder of the opportunity to protest. This explanation is unlikely in our context, however, as 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 ($p\text{-value} < 0.001$). This result demonstrates that the reminder mechanism might explain the effects of the basic aid letter only.

However, 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 in the mailing of a DCAD Notice of Appraised Value to households (see Appendix E for a sample notification). Starting on May 15, 2020, all homeowners could download their notice at the DCAD webpage. On the same day, DCAD mailed these notifications to households meeting certain criteria (e.g., increased appraised value, ownership change, loss of homestead exemption, rendered property, or new property).⁵² We mailed our letters on Wednesday, May 20, 2020. Thus, households should have received the official DCAD 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 households that did not receive the DCAD letter, since the other households would have already been reminded about the opportunity to protest through the DCAD letter. In columns (3) and (4) of Table 1, we split the results of the sample into those who were mailed or not mailed a notification from DCAD, respectively. 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 are not statistically distinguishable ($p\text{-value} = 0.317$). This finding suggests that the effect of the basic aid letter far exceeded a simple reminder effect. As a robustness check, we show that the results are robust if we conduct an even finer analysis to address potential differences between households that received the official notification or not: i.e., comparing the effect of our letters between households who were slightly to the left versus slightly to the right of the threshold for receiving the DCAD

⁵² For instance, 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%5CCEALTHALERTRecentHeadlines04032020.pdf>.

notification.⁵³

Regarding the effects of the basic aid message, our preferred interpretation is that they were mainly the product of the walkthroughs provided on the project’s website. A first piece of evidence for this interpretation relies on unsolicited feedback from participants. On the project’s website, we provided an email address to contact the researchers with concerns about the research project. Several subjects sent emails expressing gratitude for the letter and website, and many mentioned the walkthroughs. For example, some mentioned that they had wanted to protest for years but did not know how until receiving our letter. Similarly, the Field Survey included a final, open-ended question for subjects to share any thoughts with the researchers. Many used that space to express gratitude, and some explicitly stated how the information in the letter and on the website helped them navigate the protest process.⁵⁴ Another source of supporting evidence for this mechanism comes from data on our website traffic. We estimate that the basic aid message generated 903 additional direct protests.⁵⁵ We can compare this number of additional protests that were induced by the aid message to the 2,769 unique visits to the website walkthroughs, as recorded by Google Analytics (for more details, see Appendix A.4.1). Some of those visitors may have looked at the walkthroughs but did not protest, and some may have used the walkthroughs but would have protested even without them. If we assume that around one third of those website visitors were induced to protest by our website, that would explain all 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 online 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 used our suggested argument in the extra aid message “as is” to complete 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 can 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 household provided an opinion of value in their protest that is within half a percentage point of the value we suggested in their extra aid message.⁵⁶ The bottom of column (5) shows that, in

⁵³ Results presented in Appendix A.4.5.

⁵⁴ For instance, a respondent to the survey stated, “Due to your letter I will protest. In the past, I have not had this information. Your explanations and details were very helpful.”

⁵⁵ 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).

⁵⁶ 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 possibility endogeneity bias, despite the random assignment.

the control group with no letter, there was a 3.37 pp chance that a household would enter an opinion of value that coincided almost exactly with the value that we would have suggested if they were assigned to the extra aid message treatment. In other words, it is highly unlikely 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 remained equally low, which is expected because the basic aid message did not include any information about our suggested value for their household. For households that received the extra aid letter, the frequency increased by a whopping 15.287 pp (p-value<0.001). This evidence suggests that a substantial fraction of households that received the extra aid letter copied the suggested value directly into their protest forms.

There are two alternative interpretations for the effects of the extra aid message that deserve some attention. One alternative interpretation, which is consistent and is embraced by our definition of filing frictions because our letter would be providing information that is unknown to homeowners, is that households reacted to the extra aid message because they learned about the probability of a successful protest. Another alternative interpretation, which is not consistent with our definition of filing frictions, is that households reacted to the extra aid message due to fairness considerations. In Appendix A.4.4, we provide direct evidence ruling out that interpretation: using the data from the field survey, we show that the extra aid message did not affect households’ feelings of unfairness.

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 their letter or received a letter but did not read it. To correct for these types of non-compliance, we estimate the reading rate (i.e., the share of recipients that actually read the letter 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 are 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.⁵⁷ If we combine these two estimates, we arrive at a reading rate of 70.3% ($= 0.95 \cdot 0.74$). To account for this source of attenuation bias, we scale up the coefficients by a factor of 1.42 ($= \frac{1}{0.703}$). The resulting scaled-up effects would be 2.55 pp for

⁵⁷ 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.

the basic aid letter and 4.98 pp for the extra aid letter, which are still conservative scale-up factors, as there may be additional sources of non-compliance.⁵⁸

Another potential source of non-compliance is spillovers. 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 0.

To translate the filing frictions into a money-metric, we combine the results from the field experiment with the results on the homestead cap discussed in Section 3. We focus on the effect of the most comprehensive 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 filing frictions completely. For example, subjects still had to follow the instructions to file the form and may 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 using the homestead cap quasi-experiment reported in section 3.2, each \$100 reduction in the tax amount due to the homestead cap decreases the protest probability by 2.14 pp. These results indicate that the homestead cap would need to reduce the tax amount by \$232 to generate a reduction of 4.98 pp in the protest rate. These results imply that the average filing frictions cost is on the order of \$232. This is just a rough approximation. Among other things, it assumes that households care only about the costs and benefits this year, but in reality there may be dynamic considerations, too. It also assumes that the marginal households affected by the filing frictions and expected tax savings are similar.⁵⁹

As a sanity check, we compare our estimate of filing frictions to the fees charged by agents that protest on households' behalf. We identify one such company that offers the service for a flat fee. Assuming that the marginal client of this firm is indifferent between hiring this agent or protesting directly, the flat fee should constitute a measure of the filing frictions of the marginal client.⁶⁰ In 2020, the flat fee was \$139 for properties assessed below \$200,000, and \$305 for properties assessed between \$200,000 and \$500,000.⁶¹ Those flat fees (\$139

⁵⁸ For example, some households may have opened the letter too late, either after they filed a protest or after the protest deadline, whichever came first.

⁵⁹ The estimates in this section use a subset of the sample used in section 3.2. However, we show in Figure A.7 that the estimates of the effects of the expected tax savings on protests are very similar when using the subset of data we use in this section.

⁶⁰ 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 under the belief that that the agent can negotiate higher tax savings.

⁶¹ Source: https://www.dfwtaxadvisor.com/practice_areas/property-tax-protest/, accessed on May 15, 2020.

and \$305) are on the same order of magnitude as our estimated average filing frictions cost (\$232), thus suggesting that our estimates are on the right order of magnitude.⁶²

The magnitude of the filing frictions can have some policy implications too: i.e., these frictions may explain why some households (e.g., older households with potentially lower Internet literacy) protest less than others. Our findings can also explain an otherwise puzzling fact about tax protests: even though there is no fee to file a protest and no risk of increasing taxes, only a small minority of households choose to protest their taxes each year. One potential explanation is that most households do not protest because they do not want to free-ride on the taxes paid by others. Our findings point to a different explanation: most households may want to free-ride, but they expect their private costs of protesting to exceed the expected tax savings. Specifically, on the one hand, among those who protested in 2020, the average tax savings was \$338 (including both successful and unsuccessful protests). On the other hand, we estimate a cost of \$232 from filing a protest (and this is a lower-bound). The fact that the costs and benefits are on the same order of magnitude favors the interpretation that most households choose not to protest just because the private costs are too large relative to the expected savings.

5 Conditional Cooperation

5.1 Conceptual Framework

Fairness considerations can come in all shapes and sizes.⁶³ We explore one specific fairness channel that we hypothesize could be most relevant for the context of tax appeals: conditional cooperation. Robust evidence from laboratory experiments documents that conditional cooperation is significant: i.e., despite their incentives to free ride, individuals are more willing to contribute to the provision of a public good if they believe other individuals contributed too (Gächter, 2007). However, it is not obvious whether the conditional cooperation channel would be significant in a real-world context such as property taxes. Among other things, the stakes are orders of magnitude higher for property taxes than for laboratory games, and taxpayers may have strong views about the government that do not manifest in the laboratory setting (Huet-Vaughn et al., 2019).

As a preliminary step, we provide some suggestive evidence that fairness considerations

⁶² In theory, fees should not be lower than the filing frictions costs, otherwise homeowners would presumably be unlikely to protest directly. This exercise is only presented as a sanity check of the order of magnitude and, among other things, does not take into account heterogeneity across homeowners.

⁶³ Anthropologists often consider cooperation and reciprocity as features that are present among all people, or “human universals” (Brown, 1991). Economists, too, have studied fairness considerations (see e.g., Andreoni, 1995; Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Fehr and Schurtenberger, 2018).

can play an important role in households’ decisions to protest. One question from the Field Survey elicits whether the respondent feels that his or her own taxes are unfair: “Relative to the other households in the county, do you think your household pays a fair amount in property taxes?” The responses are on a scale from 1 (very fair) to 10 (very unfair). By merging the survey responses to the administrative records, we can measure whether the feeling of unfairness is correlated to the probability of filing a tax protest. The results are presented in Figure 5.a. There is a positive, strong and statistically significant ($p\text{-value} < 0.001$) association between the perceived unfairness and the probability of filing a tax appeal. Increasing the perceptions from very fair (1) to very unfair (10) is associated with a 21.1 pp increase in the probability of protesting.

As additional suggestive evidence, we employ a different question from the Field Survey that elicited households’ preferences over the distribution of tax rates. Respondents were asked to distribute a total property tax burden of \$10,000 between two hypothetical households: Household A (whose home is worth \$100,000) and Household B (whose home is worth \$400,000). The question provides a menu of seven options to choose from: Household A pays \$0 (and thus Household B pays \$10,000), \$1,000, \$2,000, \$5,000, \$8,000, \$9,000 or \$10,000. The household is asked to report whichever allocation would be most fair. If the respondent thinks that the fair allocation is for everyone to pay the same tax *rate*, then he or she would choose the third option (household A pays \$2,000 and household B pays \$8,000) in which both households pay a tax rate of 2%. If a respondent thinks that the fair thing is for everyone to pay the same tax *amount*, then she or he should pick the middle option in which both households pay \$5,000 (implying a tax rate of 5% for Household A and a tax rate of 1.25% for Household B). The responses to this question are presented in Figure 5.b. A strong majority (76.5%) of households chose the option that equalizes tax rates. In comparison, only a minority of respondents (5.47%) chose the option that equalizes the tax amounts. This evidence suggests that for a strong majority of households, the “fair” thing is for everyone to pay the same tax *rate*.⁶⁴ Thus, our conditional cooperation hypothesis is that households will be less likely to file a protest (i.e., more willing to tolerate taxes) the higher the perceived tax rate paid by the average household.

5.2 Design of the Field Experiment

One way of measuring conditional cooperation would be to randomize how much others pay in taxes and measure whether these randomly assigned average tax rates affect individuals’ decisions to protest. However, randomizing the average tax rate is not feasible. Instead of

⁶⁴ We included similar questions in the Mturk Survey (see Section 5.4) and the results were largely consistent.

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 is the second treatment arm of the field experiment introduced in Section 4. This treatment arm randomizes the content of the table we included in the middle of the first page of the letter. In Figure 3, this table is highlighted inside a red box with dashed lines (the box is for explanatory purposes and was not included in the actual letters sent to subjects). All letters include a table, but we randomize (with 50% probability) whether the table adds a column showing figures for the average Dallas home (positioned in the last column). This column shows whether the recipient’s tax rates are above or below the average Dallas home and by how much. Additionally, we cross-randomize whether the table includes a last row that makes tax rates more salient and explicit. As stated in the RCT registration, this additional treatment arm is intended as a minor robustness check and thus is discussed in the Appendix.⁶⁵

The random inclusion of an additional column is meant to provide a shock to households’ perceptions of the tax rate that 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 in response to the information shock (Cruces et al., 2013; Cullen and Perez-Truglia, 2022; 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, whereas the other half was originally upward-biased and updated in the opposite direction. In this 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 each other out. To overcome this limitation, and as anticipated in the registration, we designed the experiment based on the disclosure-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

⁶⁵ The randomization of the last row aims to address the concern that households compare tax rates (instead of tax amounts) by making them salient. For more details on the design of this treatment arm, see Appendix A.3. For the results from this treatment arm, see Appendix A.5.4. We find similar results regardless of whether the tax rate was made salient.

decision to file a protest. Additionally, we use two questions from the Field Survey as outcome variables. Appendix F provides a copy of the full survey instrument. The first question asks whether it is likely that the individual will file a protest this year. The goal of this question is to provide a validation exercise, by measuring the intention to protest as soon as the information is provided. The second and most important question, introduced above, elicits the respondent’s feeling of unfairness about his or her own property taxes. This question is designed to test the causal mechanism at play. The hypothesis is that, according to conditional cooperation, finding out that the average tax rate is higher should reduce the 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 $\bar{\tau}$ be the *actual* average tax rate in the county. 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:

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

As before, X_i^{pre} corresponds to the vector of pre-treatment control variables, which contains the same variables listed in Section 4.4. Again, this is an experiment, so the goal of using pre-treatment controls is to gain statistical power by reducing the variance of the error term (McKenzie, 2012). We 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., $(\bar{\tau} - \tau_i)$) when the true average tax rate is not disclosed. 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 average. We assume, and validate in the next section, 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 registration as the field experiment and was conducted through Amazon Mechanical Turk around the same dates as the field experiment, from June 5, 2020, to June 15, 2020. We followed several best practices for recruiting individuals in MTurk.⁶⁶ The full survey instrument is attached as Appendix G and is summarized here. This survey experiment starts by eliciting the (prior) beliefs about the average property tax rate in the subject’s county. Then, in the information-provision stage, a random half of the subjects receive accurate information about the average tax rate in their county. Next, we re-elicite the (posterior) beliefs. With this information, we 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.⁶⁷ Figure 6 presents the main results from the MTurk Survey. Figure 6.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 corresponds with accurate perceptions. Instead, the coefficient (0.237) falls significantly short of 1, indicating significant misperceptions. More precisely, Figure 6.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 expect individuals toward the left side of the x-axis in Figure 6.a to update their beliefs downwards and individuals toward the right side of the x-axis to update their beliefs upwards.

Figure 6.b corresponds to the belief updating. This figure illustrates the intuition behind the identification strategy. We 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 right of the x-axis. The x-axis in Figure 6.b is the same as in 6.a,

⁶⁶ For more details about the design and implementation of this survey, see Appendix A.5.2.

⁶⁷ Appendix A.1.3 provides descriptive statistics for this sample. We show that, in terms of observable characteristics, the MTurk sample is not identical to the other samples used in this paper but not wildly different either.

but the y-axis in 6.b corresponds to the subjects’ posterior beliefs (after the information-provision experiment) instead of their prior beliefs (as in 6.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 overestimated the average tax rate updated downwards, and individuals who underestimated it updated upwards. For a more formal test, we 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).

Figure 6 details the effects of the information shocks. It is also convenient to summarize the results in a single parameter, which we do in the econometric model (equation 7) in Section 5.3. Table 2 presents the results from this regression specification. All columns in this table are based on the same regression specification but use different samples and dependent variables. In column (1), the data are 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 $D_i \cdot (\bar{\tau} - \tau_i)$). 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.⁶⁸ This rate of pass-through is significantly above 0 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 consistent with standard learning models.⁶⁹ Indeed, this coefficient of 0.393 is similar to pass-through rates found in other survey experiments. For example, Bottan and Perez-Truglia (2020a) uses a similar research design but in the context of home price expectations and finds a pass-through rate of information shocks of 0.205.⁷⁰

⁶⁸ In Appendix A.5.3 we provide some additional robustness checks. For example, we show that the information provision on the average tax rate did not affect the respondents’ perceptions about their own tax rates.

⁶⁹ 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.

⁷⁰ This result corresponds to the coefficient reported in column (1) of Table 2 from Bottan and Perez-Truglia (2020a).

5.5 Results from the Field Experiment

Turning 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 ($\bar{\tau}$) is 0.1 pp higher than its own tax rate has a decreased 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 the dependent variables are protest indicators for the years 2015 through 2019, instead of 2020. Our letters had not been sent yet, so the information shocks should have no effect on protests in prior years. As expected, the coefficients for the other dependent variables are close to 0, 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 letters may not have been opened or opened too late. Additionally, recipients who opened letters 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 households must have read the letter to know the survey link and code needed to fill out the Field Survey. 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 column (2) shows data for all subjects letter recipients, column (3) is restricted to the subsample of 1,888 households that responded to the Field Survey.⁷¹ The survey respondents are not a random sample. In terms of household characteristics, such as home value, number of bedrooms, and tax rate, the differences between survey respondents and non-respondents are statistically significant but small.⁷² However, one substantial difference between the samples in columns (2) and (3) of Table 2 is the share of subjects who protested in 2020 (50.26% in column (3)), which is much higher than the corresponding share among subjects who received

⁷¹ The implied response rate to the survey, 3.7%, is on the same order of magnitude as the response rate in studies that sent a survey link through letters (4.7%, as reported in Sinclair et al., 2012), and substantially higher than in some studies (e.g. Perez-Truglia and Troiano, 2018; Bottan and Perez-Truglia, 2020a).

⁷² Results reported in Appendix A.5.5.

a letter (11.29% in column (2)). A natural interpretation is that the subjects who paid the most attention to our letter were those who were undecided about protesting in 2020.⁷³

The coefficient on *Information Shock* is negative (-15.392) for 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: because 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 previously discussed, the stronger effects are probably due in great part to the fact that survey respondents paid close 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. In this analysis, we estimate the same regression as in column (3) of Table 2 but the dependent variables are indicator variables for whether the respondent protested in each year during 2015–2020. That is, Figure 2.d is identical to Figure 2.c, except that the former 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 0, statistically insignificant, and precisely estimated.

Now that we focus on respondents to the Field Survey, so 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) to protest in 2020 and takes the value 0 if the household states it is unlikely (or very unlikely) to protest. Based on the evidence that the information shock affected the actual decision to protest (column (3)), we 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.5.4, we present additional robustness checks. For example, we show that the effects of the information shock are similar regardless of whether the tax rate was made more salient or not. We also show that the information provision on the average tax rate did

⁷³ Additionally, it is possible that recipients who found our letter helpful wanted to reciprocate by responding to our survey.

not affect respondents' perceptions about their own tax rates. Additionally, we show that the effect of the information shock is consequential for the number of protests and for the households' subsequent market values and tax amounts.

5.6 Causal Mechanisms of Conditional Cooperation

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

Our favorite interpretation of households' reactions to this information shock is that households changed their perception of the average tax rate and (subsequently) their feelings of unfairness. We included a question in the Field Survey to directly probe how the perceived unfairness of property taxes affects tax protest decisions. We used this question before to show that the sentiment of unfairness has a positive and statistically significant relationship with the protest probability (Figure 5.a). We now argue that under our favorite mechanism we would expect that an information shock that increases the perceived average tax rate would reduce perceptions of unfairness concerning how much a household pays in taxes. The results are presented in column (5) of Table 2, which is identical to column (3) except that the dependent variable measures the perceived unfairness on a scale from 1 (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. For example, an information shock of 0.1 pp causes a change in perceived unfairness that is equivalent to 2.14% of the standard deviation of this outcome. In comparison, the results from column (3) indicate that an information shock of 0.1 pp causes an increase in the protest probability equivalent to 3.08% of the standard deviation of that outcome.⁷⁴

Though conditional cooperation is our preferred interpretation of the effects of the information shock, and this interpretation is supported by the direct evidence from our surveys, we discuss an alternative mechanism: that subjects reacted to the information on the average tax rate because they learned from that information whether their own protests would be successful. However, below we provide evidence against this alternative mechanism.

We note that if we assume that households process information rationally, we should rule out this alternative mechanism from the beginning. A household that receives information

⁷⁴In Appendix A.5.3, we provide consistent evidence from the MTurk Survey. In that survey, 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 (holding constant the tax rates of everyone else). Consistent with the effects reported for the field experiment, we find that the information shocks from the MTurk Survey affected these fairness outcomes in the expected direction.

indicating that the average household is paying less in taxes cannot rationally infer that the reason for this fact is that its proposed value is too high. This is because differences in tax rates between households are primarily driven by different jurisdictional tax rates, exemptions, and the homestead cap. Due to non-linearities created by exemptions and the homestead cap, a successful protest can increase or decrease the tax rate a household pays. Specifically, the tax rate is computed by dividing the tax amount by the proposed value of the property, as mentioned previously. A successful protest can reduce the value of both the numerator and the denominator, and this typically leaves the tax rate roughly unchanged.⁷⁵ So, if households learn that the average tax rate is 1 pp higher than their own, it would be irrational for them to infer anything about their odds of a successful protest (for a more detailed discussion, see Appendix A.5.1).

It is possible, however, that households process information irrationally and use the information on the average tax rate to make inferences about the probability that their protest will be successful. This alternative mechanism would have heterogeneous effects on households who have protested before versus those who have not, since presumably households who protested before would have a greater understanding of how property taxes work and have more information about the odds of a successful protest from their past experience. In column (6) of Table 2, we include subjects who never protested before and in column (7) we include subjects who protested before. We find no evidence of the type of heterogeneity predicted by this alternative mechanism: the coefficients from columns (6) and (7) are similar in magnitude (-14.27 vs. -16.76) and their difference is statistically insignificant (p-value=0.83).

In addition, Appendix A.5.1 presents evidence that the odds of a successful protest is not associated with a household's proposed tax rate. In sum, though we cannot rule out that the alternative mechanism might play some role, it is very unlikely to fully explain the effects of the information shocks.

5.7 Magnitude of the Effects

One challenge for assessing the magnitude of the effect of conditional cooperation is the need to account for two forms of non-compliance. The first form of non-compliance is that some recipients may not have read the letter. As previously explained, we address this form of non-compliance by focusing on the results from column (3) of Table 2, which is the sample of recipients we are confident read the letter. However, a second form of non-compliance remains: even if they read the letter, they may not fully incorporate the feedback into their beliefs. We use the results from the MTurk Survey to correct for this additional form

⁷⁵ A successful protest leads to a marginal *increase* in the tax rate if a household has a binding homestead cap.

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 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 recipients' perception of the average tax rate paid in the county by 0.1 pp would decrease their protest probability by 3.90 pp ($= 0.1 \cdot 15.392 \cdot 2.54$). Thus, subjects are responsive to their perceptions of taxes paid by everyone else. Moreover, conditional cooperation is just one possible manifestation of fairness considerations, so this provides a lower bound for the importance of fairness considerations.⁷⁶

Last, combining our estimates allows us to provide a back of the envelope estimate of the willingness to pay for fairness, in dollar terms. The results from this section indicate that a 0.1 pp increase in the perceived average tax rate would decrease the protest probability by 3.90 pp. The results from Section 3 above indicate that each \$100 increase in the marginal benefits from protesting causes a 2.14 pp reduction in the probability of protesting. Combining these two findings, we estimate that a 0.1 pp increase in the perceived average tax rate increases the willingness to pay taxes by about \$182.24 ($= \frac{3.9 \cdot 100}{2.14}$).

6 Heterogeneity by Political Party

Recent evidence from surveys and laboratory experiments suggests that differences in beliefs and preferences across partisan lines may be important in how taxpayers interact with the government (e.g., Huet-Vaughn et al., 2019). Motivated by those findings, we explore whether partisanship matters for the decision to file a protest.

6.1 Data on Party Affiliation

To split the analysis by Republican and Democratic 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, individuals 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 Republican primaries would be classified as a 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, it is public record when an individual contributes over \$200 to a Democratic or Republican candidate (Perez-Truglia and Cruces,

⁷⁶ For example, in the spirit of benefit-based taxation (Weinzierl, 2018), households whose kids do not attend public school may protest because they consider it unfair that most of the tax revenues go to public schools.

2017). 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.⁷⁷

We classified each of the 423,607 single-family homes in the main sample as more likely to identify as Republican or Democrat.⁷⁸ Specifically, we identified 57% of subjects as more likely to be Democrats and the remaining 43% as more likely to be Republican. Indeed, this narrow lead in support of the Democratic party is consistent with recent electoral results from Dallas County. In the 2012 presidential election, Barack Obama received 57% of the votes, whereas Mitt Romney received 42%, and the remaining 1% of votes went to third-party candidates.

6.2 Protest Rates by Party Affiliation

To the extent that Republicans indicate they support taxation less than Democrats (Stantcheva, 2020), one might expect Republicans to be much more likely to protest their property taxes than Democrats. These differences are explored in the top half of Figure 7, which shows the protest probabilities (combining direct protests and protests through agents) for Republicans and Democrats.⁷⁹ These protest rates are calculated for different groups of home values: the first group corresponds to properties valued under \$100,000, while the last group corresponds to properties valued above \$500,000. The bottom half of the figure shows the percentage of Democratic and Republican homes in each group. Relative to Democrats, Republicans tend to live in more expensive homes. If owners of more expensive homes are more likely to protest, that could mechanically generate differences in protest rates across partisan lines. We find that, on average, Republicans were more likely to protest than Democrats: in the sample of 423,607 households, 21.98% of Republicans protested in 2020, compared to 13.41% of Democrats. However, this difference is due to differences in home values: Figure 7 shows that, when comparing homes of roughly similar value, the differences in protest rates between Democrats and Republicans are small. For example, for homes in the median category, valued between \$200K–\$300K, the protest probabilities are 14.23 pp for Republicans versus 13.50 pp for Democrats. This difference is statistically significant (p-value=0.058) but small in magnitude. For the other groups, the differences are sometimes statistically significant,

⁷⁷ Results presented in Appendix A.6.1.

⁷⁸ For individuals with missing information, such as those not registered to vote, we use a simple predictive model to impute their party affiliation. For details on this imputation, see Appendix A.6.1.

⁷⁹ The results are broadly similar if we focus on direct protests only (results reported in Appendix A.6.2). Specifically, Democrats protest directly slightly more frequently than Republicans and Republicans use agents slightly more often than Democrats.

but always economically small.

6.3 Motives for Protesting

Even if Democrats and Republicans protest with a similar probability when accounting for differences in wealth, it is possible that they protest for different motives. In Sections 3–5, we analyzed the roles of selfish (i.e., expected tax savings and filing frictions) and fairness motives. Here we break that analysis down by political party.

We start with the response to expected tax savings. Figure 8 reproduces the results from Figures 1.a and 1.c on the effects of the homestead cap and compares Democrats (left panels: (a) and (c)) and Republicans (right panels: (b) and (d)). Exceeding the homestead cap threshold by \$10,000 causes a decrease in the protest rate of 4.25 pp (p-value<0.001) among Democrats (from Figure 8.c) and 4.70 pp (p-value<0.001) among Republicans (from Figure 8.d). These effects are close to each other and statistically indistinguishable (p-value=0.597), suggesting similar responsiveness to the expected tax savings among Republicans and Democrats.

The differences remain small if we instead normalize the effects of the homestead cap on protest rates by the corresponding effects on tax amounts (Figures 8.a and 8.b). Exceeding the homestead cap threshold by \$10,000 causes a decrease in the tax amount of \$220 (p-value<0.001) among Democrats (from Figure 8.a) and \$204 (p-value<0.001) among Republicans (from Figure 8.b), and the difference between these two coefficients is small and statistically insignificant (p-value=0.674). In other words, for Democrats, a \$100 reduction in the tax amount caused by the homestead cap is associated with a 2.03 pp drop in the protest rate, whereas the corresponding effect is 2.31 pp for Republicans. This evidence suggests that both Republicans and Democrats are highly responsive to the expected tax savings of protesting, but Republicans seem, if anything, to be slightly more responsive.

Next, we look at the partisan differences in responses to filing frictions and show the results in the last two columns of Table 1. These columns reproduce the results from the baseline specification of column (1) of Table 1 and compare subsamples of Democrats (column (6)) and Republicans (column (7)). As with the response to expected tax savings, we find the response to filing frictions 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 the one hand, the coefficient on the 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)), and the difference is statistically significant (p-value=0.065). Together, the evidence suggests that both Republicans and Democrats are elastic to the filing frictions of protesting, but Republicans may be somewhat more elastic.

The evidence on the role of filing frictions and expected tax savings 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. Although fairness considerations are universal, results from laboratory experiments suggest large differences across individuals in the strength of conditional cooperation.⁸⁰ 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 Democratic 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)).⁸¹ These point estimates suggest that conditional cooperation is somewhat stronger among Democrats than among Republicans, but the difference is statistically insignificant (p-value=0.515). In sum, the evidence cannot rule out that conditional cooperation is equally important for Democrats and Republicans.

7 Conclusions

Individuals in all U.S. states can legally file a protest to reduce their property taxes. Using experimental and quasi-experimental methods, we provide evidence on the determinants of the decision to file a protest. We show that expected tax savings, filing frictions, and fairness considerations play a major role, while partisanship does not seem to matter at all.

Although our evidence is based on data from a specific U.S. county, to the extent that tax protests work similarly across counties both within and outside Texas, the results should be generalizable to those other settings. Indeed, conducting the same experimental and quasi-experimental designs from our study in other U.S. counties would be straightforward. For instance, other counties also have a homestead cap for property taxes, and our mailing campaign could be readily conducted in many other counties. In this spirit, we provide detailed accounts of the implementation and data sources that other researchers can follow, and we are happy to share data, code, tips, and additional resources. Moreover, we believe

⁸⁰ For example, 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).

⁸¹ Appendix A.6.3 presents an event-study falsification test for these results.

that our framework can be adapted to study research questions beyond tax compliance and in diverse fields such as political economy, public economics, finance, and behavioral economics.

Our novel setting has several features that we believe 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 for non-disclosure agreements or data user 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. Last, the experiment can be implemented on massive scales, involving potentially up to millions of subjects.⁸²

We conclude by summarizing some policy implications. Our findings uncover some aspects of the protest system that deserve further attention and research. For example, large differences in protest rates occur across wealthier and less wealthy households. Further, even after differences in home values are accounted for, some groups, such as ethnic minorities, are substantially less likely to file a protest. The large filing frictions that we document in this paper suggest that, as claimed by critics (Lieber, 2020), the demographic differences in protest rates may reflect inequitable access to the system of tax appeals. Indeed, our findings hint at some low-cost interventions that can be used to mitigate the inequities in the system. For example, Hispanic households are 3.61 pp less likely to protest than comparable White households.⁸³ We find that our letter with extra aid increases the protest rate among Hispanic households by 2.55 pp (p-value<0.001).⁸⁴ Thus, a low-cost mailing intervention like ours targeted towards poorer and/or Hispanic households could go a long way into reducing the inequity in access.⁸⁵ These low-cost interventions could be promoted either internally by county assessors' offices or externally through nongovernmental organizations. For example, a nongovernmental organization has been recently registered in Texas to help with the property tax appeals of poor and otherwise disadvantaged households (Property Tax Appeals Justice, NGO). These interventions could be targeted not only to poorer or Hispanic households but to any other groups left behind by the system of tax appeals.

⁸² In Dallas County alone, it can potentially involve hundreds of thousands of subjects. By pooling multiple counties, it could be scalable to millions of subjects.

⁸³ For details, see Appendix A.1.4.

⁸⁴ This result corresponds to the regression from column (1) of Table 1, but estimated with the 20.39% of the sample that was classified as Hispanic.

⁸⁵ For related examples on targeted interventions of this nature, see for example Thaler and Sunstein (2009) and Finkelstein and Notowidigdo (2019).

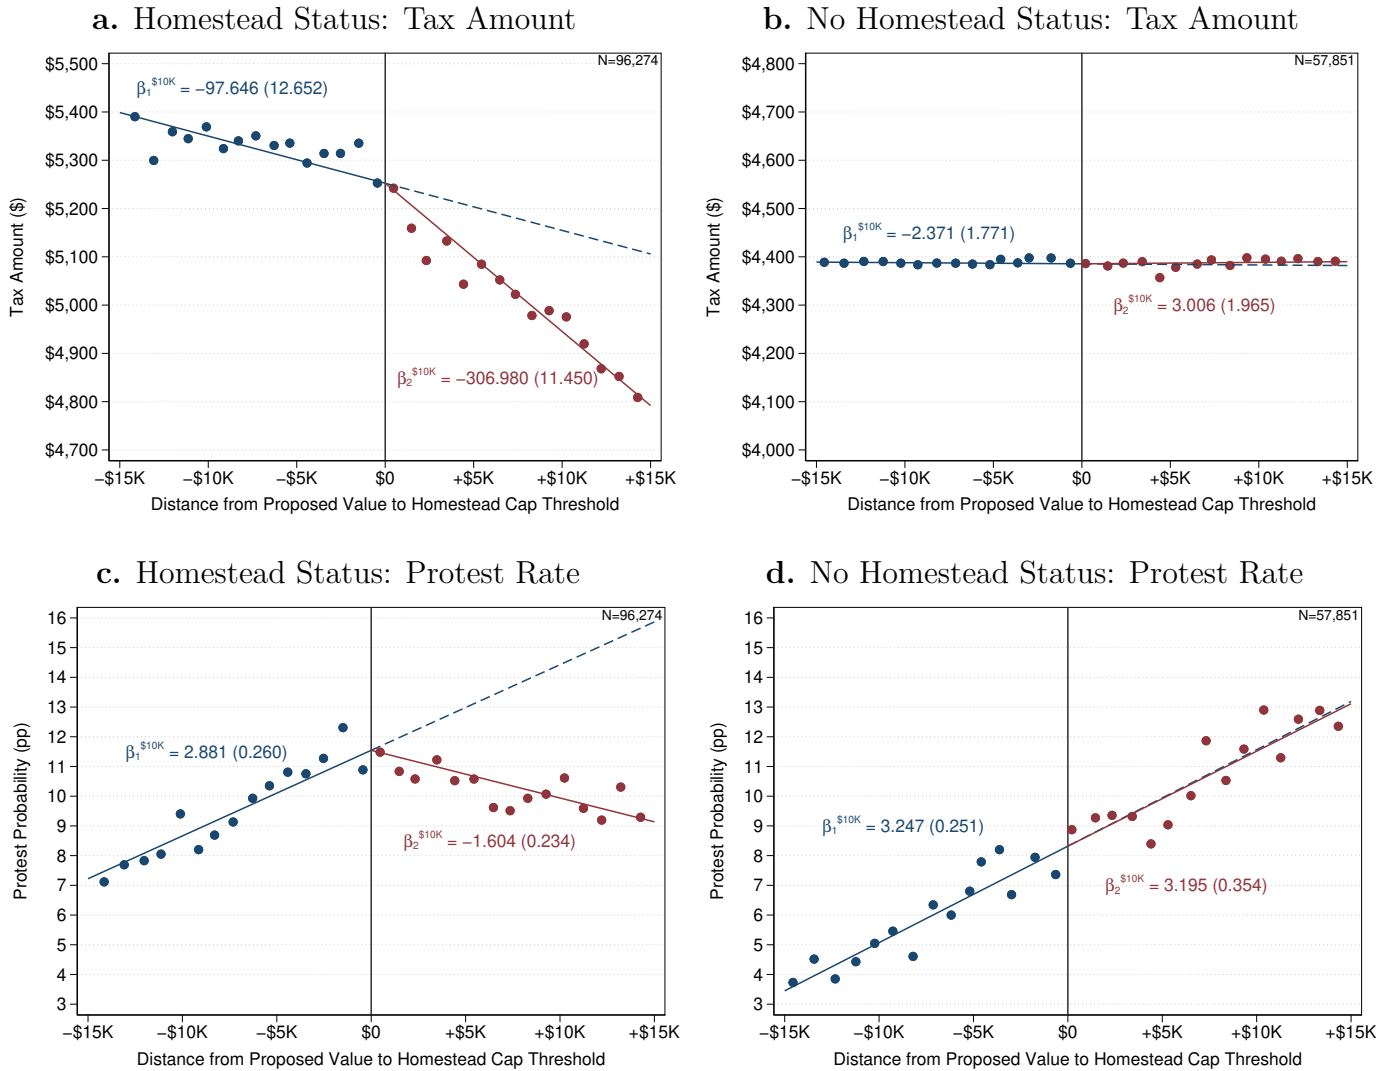
References

- Ahler, D. J. (2014). Self-Fulfilling Misperceptions of Public Polarization. *The Journal of Politics* 76(3), 607–620.
- Alatas, V., A. Banerjee, R. Hanna, B. A. Olken, R. Purnamasari, and M. Wai-Poi (2016). Self-Targeting: Evidence from a Field Experiment in Indonesia. *Journal of Political Economy* 124(2), 371–427.
- Alatas, V., A. Banerjee, R. Hanna, B. A. Olken, and J. Tobias (2012, 6). Targeting the Poor: Evidence from a Field Experiment in Indonesia. *American Economic Review* 102(4), 1206–1240.
- 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*.
- 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 (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*.
- 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.
- Cruces, G., R. Perez-Truglia, and M. Tetaz (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics* 98, 100–112.
- Cullen, J., N. Turner, and E. Washington (2020). Political alignment, attitudes toward government

- and tax evasion. *American Economic Journal: Economic Policy*, forthcoming.
- Cullen, Z. and R. Perez-Truglia (2022). How Much Does Your Boss Make? The Effects of Salary Comparisons. *Journal of Political Economy* 130(3), 766–822.
- 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.
- Doerner, W. M. and K. R. Ihlanfeldt (2015). The Role of Representative Agents in the Property Tax Appeals Process. *National Tax Journal* 68(1), 59–92.
- 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.
- 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.
- Giaccobasso, M., B. Nathan, R. Perez-Truglia, and A. Zentner (2022). Where Do My Tax Dollars Go? Tax Morale Effects of Perceived Government Spending. *NBER Working Paper No. 29789*.
- Goolsbee, A. (2006). The Simple Return: Reducing America’s Tax Burden Through Return-Free Filing. *The Hamilton Project Discussion Paper No. 2006-04*.
- Hayashi, A. T. (2014). The Legal Salience of Taxation. *The University of Chicago Law Review* 81(4), 1443–1507.
- 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*.
- Lieber, D. (2020). Why protesting your property appraisal is so hard. *The Dallas Morning News*, June 4th 2020.
- 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.


- Parker, W. and N. Friedman (2021). Zillow Quits Home-Flipping Business, Cities Inability to Forecast Prices. *The Wall Street Journal*, November 2 2021.
- 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.
- Sinclair, M., J. O’Toole, M. Malawaraarachchi, and K. Leder (2012). Comparison of response rates and cost-effectiveness for a community-based survey: postal, internet and telephone modes with generic or personalised recruitment approaches. *BMC Medical Research Methodology* 12(1), 132.
- Slemrod, J. (2019). Tax Compliance and Enforcement. *Journal of Economic Literature* 57(4), 904–954.
- Stantcheva, S. (2020). Understanding Tax Policy: How Do People Reason? *NBER Working Paper No. 27699*.
- Stein, J. and C. Ingraham (2019). Corporations paid 11.3 percent tax rate last year. *The Washington Post*, December 16, 2019.
- 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.
- Urban Institute (2021). State and Local Finance Initiative. Accessed from: <https://www.urban.org/policy-centers/cross-center-initiatives/state-and-local-finance-initiative/projects/state-and-local-backgrounders/property-taxes>.
- U.S. Census Bureau (2021). Population, Dallas County, Texas. <https://www.census.gov/quickfacts/fact/table/dallascountytexas/POP010220>.
- U.S. Monitor (2014). 7 Myths of Direct Mailing. Retrieved March 28, 2020, from <https://www.targetmarketingmag.com/promo/7MythsofDM.pdf>.
- Villanueva, C. (2018). What is Recapture? *Center for Public Policy Priorities Report*, August 30, 2018.
- 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 Amount 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 the red diamonds represent the effects of the extra aid letter. Panel (c): The blue dots represent the coefficient on the information shock ($D_i \cdot (\bar{\tau} - \tau_i)$) 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
<i>Proposed Value</i>	\$174,810	\$294,846
<i>Estimated Tax Amount</i>	\$3,057	\$5,916
<i>Estimated Tax Rate</i>	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.8. 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 Ronyan 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!

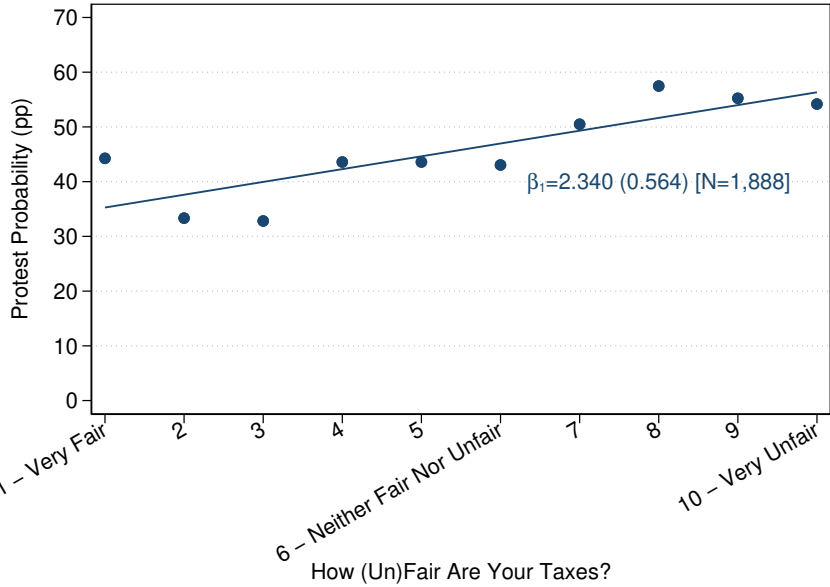
Alejandro Zentner
Associate Professor
University of Texas at Dallas

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

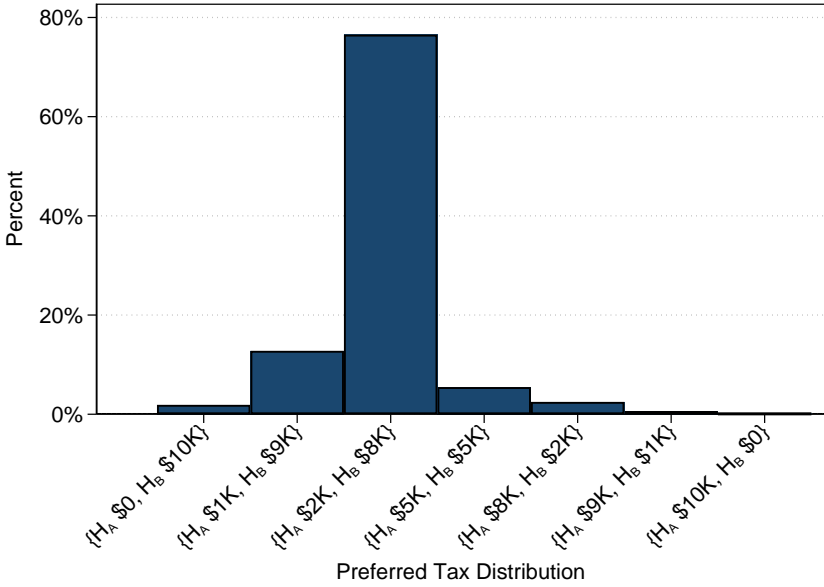
Notes: A sample of the second page of the letter used in the field experiment. Panel (a) does not contain the Extra Aid Message, while panel (b) does in the section framed with the red dashed lines. This red frame in panel (b) was added to this figure for emphasis but does not appear in the actual letters.

Figure 5: Fairness, Property Taxes, and Protests

a. Direct Protest Rates by Perceived Unfairness of Own Taxes



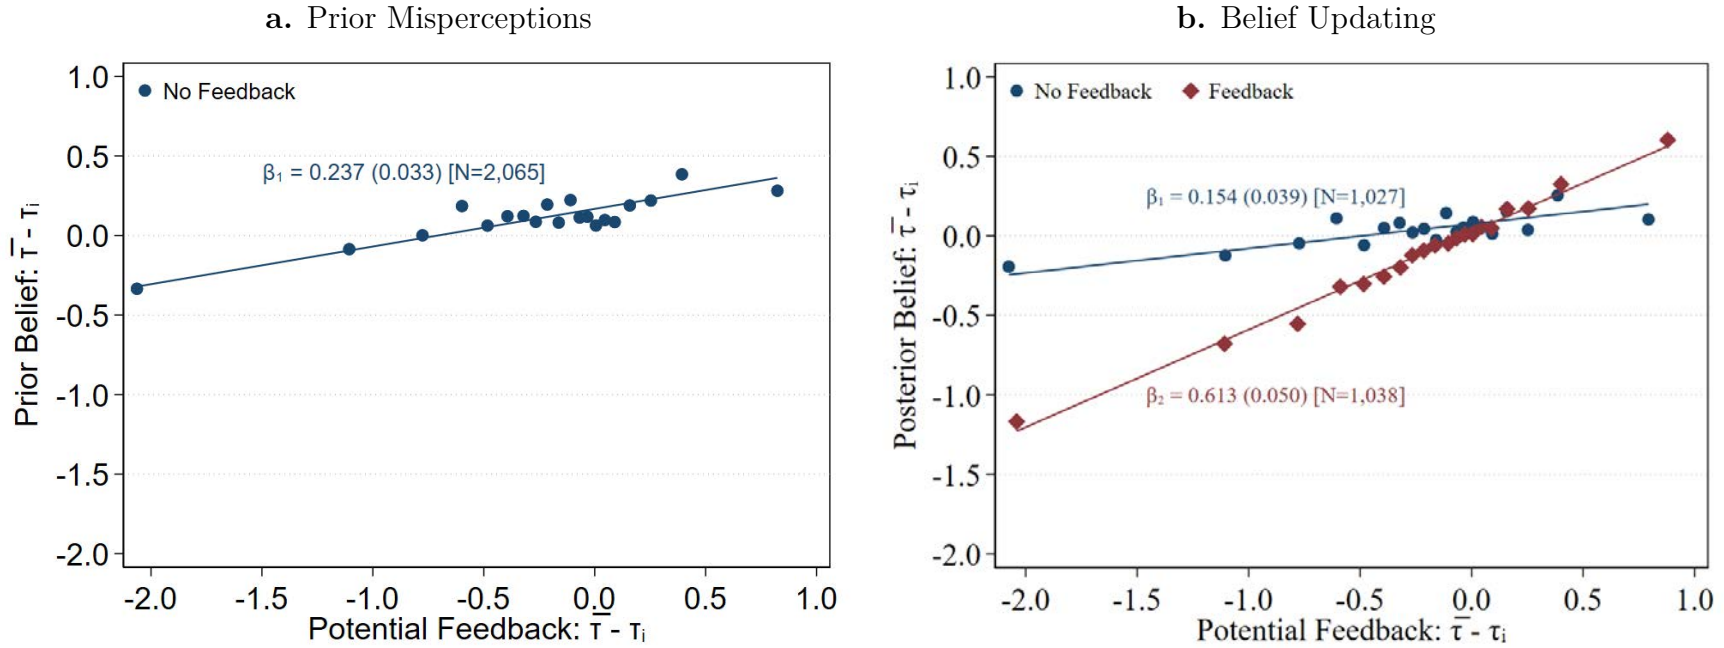
b. Social Norms about the Fair Distribution of Taxes



49

Notes: Panel (a) contains point estimates with a fitted line. This figure plots direct protest rates versus households' perceived unfairness of their own property taxes, based on a question included in the Field Survey about whether the taxes of the respondent are unfair relative to the taxes of everyone else, on a 1-10 scale. Sample includes 1,888 survey responses. Panel (b) shows the distribution of social norms based on a question included in the Field Survey regarding the fair tax burden distribution. The horizontal axis shows the options presented to respondents on how to distribute a total tax burden of \$10,000 between Household A (which is worth \$100,000) and Household B (which is worth \$400,000). The vertical axis presents the share of the responses choosing each option.

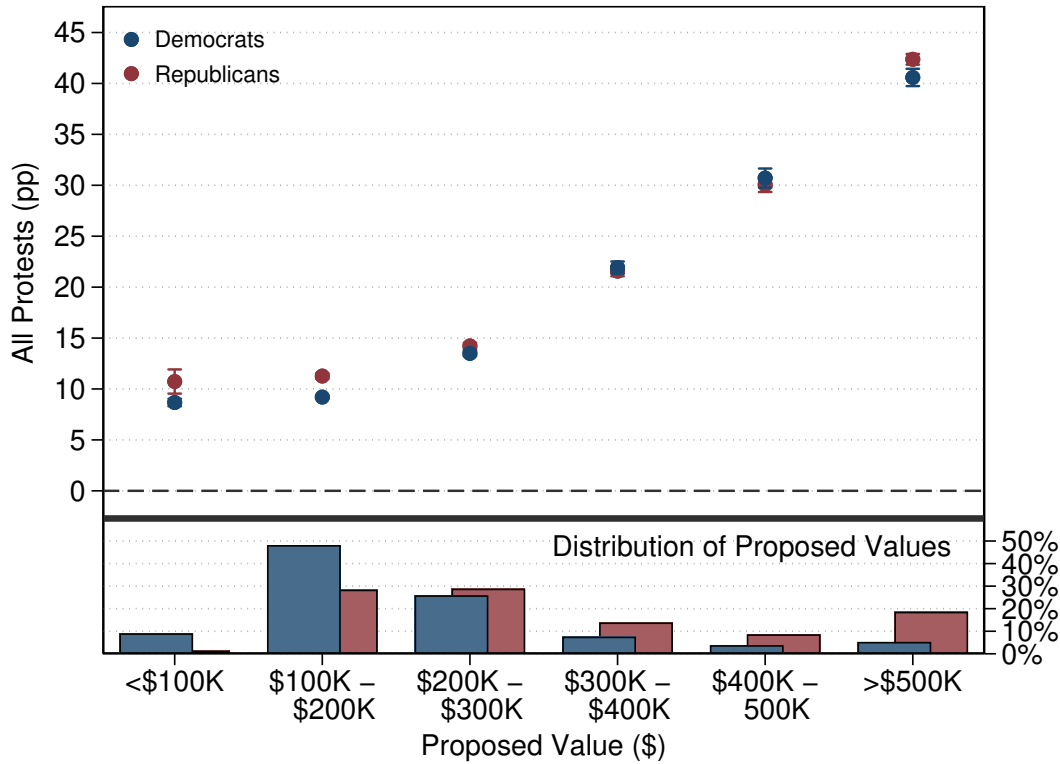
Figure 6: Prior Misperceptions and Belief Updating in the Mturk Survey



50

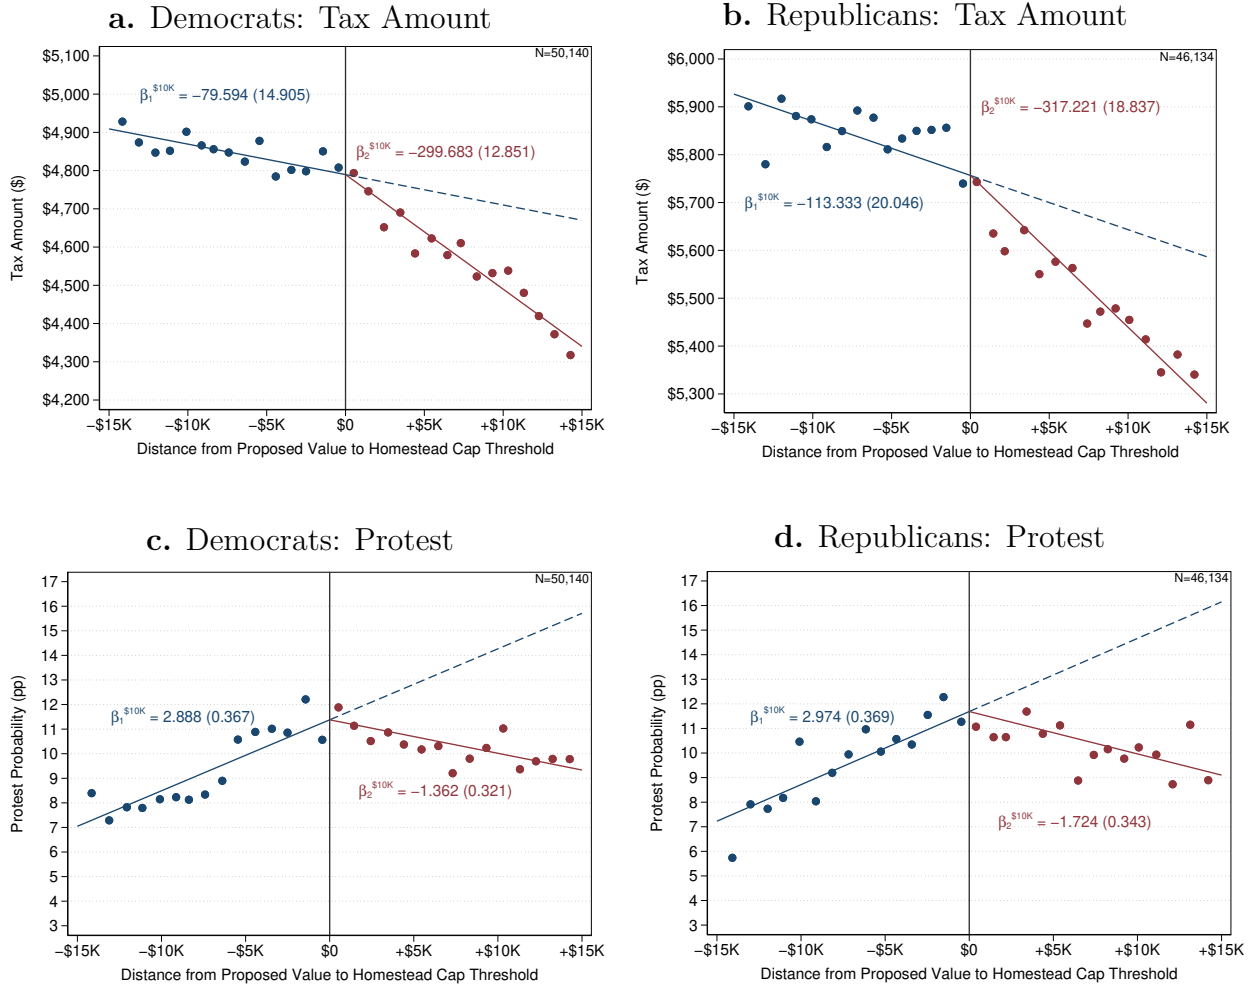
Notes: This figure shows binned scatterplots based on the Mturk Survey. Each line corresponds to a separate OLS regression, with robust standard errors in parentheses and the number of observations in brackets. $\bar{\tau} - \tau_i$ refers to the difference between the average tax rate in the respondent's 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 7: Heterogeneity in Protest Rates by Political Party



Notes: Point estimates with 95% confidence intervals in brackets in the top panel, based on robust standard errors. The sample in the top and bottom panels is the main sample of 423,607 single-family homes in 2020. This figure corresponds to all protests (i.e., direct protest as well as protest through agents). The share of households in each proposed value bracket by party is reported in the bottom panel. Households' political affiliation is imputed by merging the taxpayer records with various sources of data such as participation in the primaries and campaign contributions provided by a private vendor (Aristotle International). For individuals with missing information, we use a simple predictive model to impute their party affiliation. The proxy for political party provided by the vendor is highly consistent with voting data at the precinct-level, as discussed in Section 6.

Figure 8: 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)	(2)	(3)	(4)	(5)	(6)	(7)
	P_{2020}	P_{2020}^{won}	P_{2020}	P_{2020}	SO_{2020}	P_{2020}	P_{2020}
Basic Aid Letter ⁽ⁱ⁾	1.792*** (0.249)	1.213*** (0.222)	1.449*** (0.347)	1.935*** (0.339)	0.795 (0.719)	1.943*** (0.317)	1.509*** (0.394)
Extra Aid Letter ⁽ⁱⁱ⁾	3.509*** (0.258)	2.621*** (0.231)	3.108*** (0.364)	3.745*** (0.350)	15.287*** (0.979)	3.027*** (0.326)	3.994*** (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)
	$\bar{\tau}_i^{post}$	P_{2020}	P_{2020}	I_{2020}	U_{2020}	P_{2020}	P_{2020}	P_{2020}	P_{2020}
Information Shock ($\bar{\tau}$)	0.393*** (0.071)	-0.950* (0.509)	-15.392*** (5.591)	-13.220*** (4.575)	-0.468* (0.243)	-14.277** (7.166)	-16.764* (9.248)	-18.317** (8.204)	-10.922 (7.837)
Field Survey Subsample			✓	✓	✓	✓ I	✓ II	✓ Dem.	✓ Rep.
Mean Outcome	1.24	11.29	50.26	81.90	7.12	46.53	54.55	50.47	50.10
Std. Dev. Outcome	0.77	31.65	50.01	38.52	2.15	49.90	49.82	50.03	50.02
Observations	2,065	50,394	1,888	1,867	1,888	1,008	880	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* ($\bar{\tau}$) corresponds to the information shock term ($D_i \cdot (\bar{\tau} - \tau_i)$). 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: $\bar{\tau}_i^{post}$ 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 did not protest during 2015 through 2019 (column (6)) and ii) subjects who protested at least once during 2015 through 2019 (column (7)). Columns (8) and (9) split the sample in column (3) between likely Democrat or Republican.