NBER WORKING PAPER SERIES

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

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

We are thankful for excellent comments from colleagues. This project was reviewed and approved in advance by the Institutional Review Board at The University of Texas at Dallas. The experiments were pre-registered in the AEA RCT Registry (#0005992). Adrian Cadena Medina, Luisa Cefala, Dongwook Chun, Karl Dill and Santiago De Martini provided superb research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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

ABSTRACT

In all U.S. states, individuals can file a protest with the goal of legally reducing their property taxes. This choice provides a unique opportunity to study preferences for redistribution via revealed preference. We study the motives driving tax protests through two sources of causal identification: a quasi-experiment and a pre-registered large-scale natural field experiment. We show that, consistent with selfish motives, households are highly elastic to their private benefits and private costs from protesting. We also find that social preferences are a significant motive: consistent with conditional cooperation, households are willing to pay higher tax rates if they perceive that others pay high tax rates too. Lastly, we document significant differences between the motivations of Democrats and Republicans.

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

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

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

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

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

1 Introduction

Taxation and redistribution are prerequisites for the success of modern democracies, and among the oldest topics of study in political economy (Alesina and Glaeser, 2004). However, studying redistribution preferences presents significant measurement challenges and, as a result, the existing evidence is largely restricted to survey data (Alesina and Giuliano, 2011). In this paper, we exploit a novel context that provides revealed preference evidence on individuals' support for government taxation and redistribution: tax protests.

In many U.S. counties, individuals can file a protest with the goal of legally reducing their property taxes. These protests are consequential: while there is no guarantee that a protest will reduce one's tax bill, in practice they often do. In this paper, we study the choice to protest taxes, and use the data to shed light on the roots of preferences over taxation and redistribution.

Although tax protests are allowed across the country, our analysis focuses on one of the largest counties in the United States: Dallas County, Texas. We look at a specific county for practical reasons – among them that it is easier to implement a field experiment in a single county. Still, to the extent that protests work similarly across counties – and they do work similarly in at least all counties in Texas – our results should be generalizable. Moreover, studying property taxes in Texas is especially relevant, as the stakes are high: because Texas does not have a state income tax, property taxes are a key source of revenue for the provision of government services. For example, the average household from Dallas County is expected to pay around \$5,916 in property taxes in 2020, corresponding to a tax rate of 2.01% on their home's market value.[1]

The protest process in Dallas county can be summarized as follows. The Dallas Central Appraisal District (DCAD) formulates a proposed assessment of the property's market value and notifies the household of it (hereinafter, we refer to this amount as the *proposed value.* The property taxes will be calculated based on this proposed value. The household can file a protest, arguing, for example, that the proposed value is too high.[2] Next, owners can protest on their own (which is the main focus of this paper) or they can hire an agent to protest on their behalf. If the protest is successful, the effective assessed value as well as the respective tax bill will be reduced. In 2020, a total of 8.40% of households in Dallas County filed a

[1] These statistics are based on administrative data and focus on the universe of single-family homes. Unless stated otherwise, all the statistics provided in this paper are based on this sample.

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

protest of their property taxes on their own, and an additional 8.42% of households protested with the help of an agent, resulting in a total protest rate of 16.83%. We estimate that 49.4% of the protests in 2020 were successful and that these successful protests resulted, on average, in $601 in tax savings.[3]

Teasing apart the motives for protesting from the raw data is challenging. Take, for example, the fact that only a minority of households choose to protest each year. One interpretation of this fact could be that the majority of households support taxation and redistribution and would not want to free-ride on the taxes paid by others. A very different interpretation is that most households are selfish and would prefer to free-ride, but they still choose not to protest because their private cost from protesting is greater than their expected tax savings. In this paper, we use publicly available administrative data and exploit quasi-experimental and experimental variation to address these questions.

First, we exploit quasi-experimental variation to measure how private benefits from protesting affect the decision to file a protest. All counties in Texas must use a cap when calculating taxes for households with homestead status.[4] This cap generates a sharp, discontinuous kink in the expected benefits from protesting. The analysis of this sharp kink indicates that households are more likely to protest when they stand to gain more from protesting. A 0.1 percentage point (pp) increase in the tax rate cap causes an increase in the probability of protesting by 3.65 pp. This effect translates into a 7.34 elasticity – meaning that the decision to protest is highly elastic to the financial benefits.

Second, we study the role of hassle costs. Any household can protest for free, so there are no pecuniary costs from protesting; however, households may still incur hassle costs: protesting takes time, and some households may have trouble figuring out how to do it. We conducted a pre-registered field experiment to create exogenous variation in the hassle costs of protesting. We selected a subject pool of 78,462 households and sent letters to a random sample of 50,394 with helpful information on how to file a protest. As a whole, this subject pool was estimated to pay $560 million dollars in property taxes in 2020. In the letters, we randomized the intensity of our aid provision. The basic letter treatment included information such as step-by-step guides for filling out the forms by mail or online. A second letter type included extra instructions based on the fact that, when protesting property taxes, one of the most challenging aspects of the process is preparing an argument to support the

[3] These statistics are for households that protested on their own, as a significant fraction of the protests through agents have not been resolved by the time we conducted the analysis. This estimate and other estimates related to the outcomes of 2020 protests are still preliminary and will be updated with the latest administrative data in a future version of this paper.

[4] Homestead is a legal status that can be granted to properties that constitute the main residence for the household. This legal status comes with several benefits related to property taxes (e.g., exemptions and caps) and other benefits such as exemption from forced sale for collection of debt.

request. The extra aid message included an argument tailored to each recipient, who could simply copy-paste it into their own protest form. These letters presented information about a comparison property near the recipient's own property that was similar in all observable characteristics and had been recently sold for a lower price than the market value proposed by DCAD.

We find that the letters had a large impact on the probability of filing a protest; and while the basic aid was helpful, the impact was significantly higher when the letter included the additional instructions. The letter with extra aid led to a protest rate increase of 3.51 pp; for comparison, the fraction of households protesting in the control group that received no letters was 8.67 pp. In other words, a simple intervention by mail increased the protest rate by a whopping 40.5%. This evidence suggests that hassle costs are of first order in the decision to protest taxes.[5] Moreover, we combine the experimental and quasi-experimental estimates to quantify the hassle costs. Our back-of-the-envelope calculations suggest that, at the margin, the hassle cost of protesting amounts to $226. Indeed, this estimate constitutes just a lower bound, as our intervention reduces some of the hassle costs but is far from eradicating them completely.

These results have important implications for the preferences for redistribution. For households that protested on their own in 2020, the average reduction in the tax bill was $297 (including successful and unsuccessful protests). Because the great majority of households do not protest, one naive interpretation could be that the great majority of households are not selfish and are happy to pay the extra $297 in taxes; however, the fact that the hassle costs are in the order of $226 challenges this interpretation. Given the magnitude of average costs and average benefits, a more plausible interpretation is that most households would want to reduce their own taxes, but do not protest because the hassle costs are just too large relative to their expected savings.

That the evidence presented here suggests that selfishness plays a significant role in the decision to protest taxes does not mean that social preferences cannot play a role as well. We introduced an additional treatment arm in the field experiment to measure conditional cooperation: that is, whether households are more willing to tolerate a higher tax rate if they perceive the average citizen as facing a higher tax rate, too. Conditional cooperation has been documented in laboratory public good games (Gächter, 2007). However, it is not clear whether the results from the laboratory would be economically significant with real-world issues such as property taxes. Among other things, the stakes are orders of magnitude higher for property taxes than for laboratory games, and individuals may have strong views about

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

the government that do not manifest in laboratory games (Huet-Vaughn et al., 2019).

Our identification strategy to study conditional cooperation is based on an information-provision experiment embedded in the field experiment: we randomized whether the letter sent to each household included accurate information on the average tax rate in the county. To the extent that households have systematic misperceptions, the information experiment can generate exogenous variation in subjects' perceptions of the average tax rate in the county. To validate this research design, we simultaneously conducted a separate, pre-registered, survey experiment with a different sample of respondents.

Results from the field experiment suggest that the average household displays conditional cooperation: when persuaded that the average tax rate in the county is higher, they become less likely to protest their own taxes. This effect is not only statistically significant, but also economically large. To appreciate the quantitative importance of conditional cooperation, consider the following counterfactual analysis based on our preferred estimates combined with our results from the homestead cap quasi-experiment. If we increased a household's tax rate by 0.1 pp while holding everyone else's tax rates constant, the protest probability for that household would increase by 9.63 pp. If, instead, we increased the household's tax rate by 0.1 pp while increasing everyone else's tax rates by 0.1 pp, too, then the household's protest probability would still increase but by 6.43 pp. Thus, the conditional cooperation preferences dampen a third of the effect of the tax hike.

Last, we explore heterogeneity by partisan identity in the motives for protesting. Evidence from survey data suggests that Republicans and Democrats have significantly different opinions on preferences for redistribution and taxation (Stantcheva, 2020). By comparing their preferences for protesting, we can explore whether these differences exist when Republicans and Democrats face real stakes. We matched the individual-level data from the property tax records to the Texas voter files to construct a proxy for whether each taxpayer is likely to be Democrat or Republican. We show that there are some systematic differences between Democrats and Republicans regarding their propensities to protest. Although both Democrats and Republicans are highly elastic to the private costs and private benefits of protesting, there are some significant differences in magnitude: relative to Democrats, Republicans are somewhat more elastic to their private incentives. We also find that Democrats care a lot more about conditional cooperation than Republicans, but the difference is imprecisely estimated so this result must be taken with a grain of salt.[6]

This study is both related and contributes to various strands of scholarship. First, this study contributes to the literature on the decision to protest taxes. Protesting property

[6] Our evidence is consistent with the findings from Cullen et al. (2020) suggesting that partisanship can affect tax compliance.

taxes is allowed in all 50 U.S. states (Dobay et al., 2019).[7] Protesting property taxes is allowed outside of the United States, too (Dobay et al., 2019; Group, 2019).[8] Despite being so common, tax protests have received little attention in the economics literature. To the best of our knowledge, there are two exemptions: Jones (2019) uses data on the decision to protest taxes to provide a test of loss aversion,[9] and Avenancio-León and Howard (2019) show that local governments place a disproportionate fiscal burden on racial and ethnic minorities and further document that some of those differences operate through tax appeals.[10] We contribute to this literature by being the first to use tax protests to learn about preferences for redistribution. Furthermore, we identify and measure three important factors that affect the protest decision: private benefits, private costs, and conditional cooperation. Indeed, these findings can provide key inputs for the design of a more equitable system of tax appeals.

Second, we contribute to the literature on preferences for redistribution. Revealed preference evidence has proved elusive in this body of literature because individuals do not get to choose their tax rates. The typical way in which economists study preferences for redistribution is through survey data (Alesina and Giuliano, 2011; Cruces et al., 2013; Kuziemko et al., 2015), which has a number of well-known limitations – for example, individuals may say that they want high redistribution due to social pressure, but when the stakes are high they may act selfishly instead. Few efforts have been made to study preferences for redistribution with revealed-preferences data, such as using survey data on voting preferences (Fisman et al., 2017; Epper et al., 2020) or charitable giving (Fisman et al., 2007), but they all have their own sets of limitations.[11] We contribute to this literature by being the first to study preferences for redistribution using the decision to protest taxes. While this context is not ideal,

[7] In practice, protesting property taxes may be more common in some counties than in others, for example due to institutional differences. As an illustration, while none of the counties in Texas charge a fee to protest property taxes, some counties outside of Texas can charge fees, which can be substantial in some cases (Dobay et al., 2019).

[8] For example, Dobay et al. (2019) found that protesting property taxes was allowed in all 10 countries that were examined, and Group (2019) shows that property tax protests are allowed in several Latin American countries too.

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

[10] Some studies look at property taxes more generally, without focusing on protests. One recent example is Cabral and Hoxby (2012), providing evidence of how salience affects property tax rates and limits.

[11] Voting is secret, and thus can only be observed at the individual level with survey data. Furthermore, voting is inconsequential in that the probability of a single vote being pivotal is often negligible. While preferences for redistribution is certainly one issue that people have in mind when deciding whom to vote for, the decision can depend largely on non-economic factors such as abortion rights and Second Amendment rights. Charitable giving constitutes a much smaller fraction of GDP. Giving is largely driven by factors other than preferences for redistribution, such as religious participation (Bottan and Perez-Truglia, 2015). Moreover, giving removes the government from the equation, which can be misleading: some people may be willing to give to charity but would still not want government redistribution.

we believe it provides significant improvement in a number of respects. Most notably, tax protesting is a type of high stakes behavior that can be measured objectively with administrative records. Indeed, we provide detailed instructions so that other researchers can use this same experimental framework to study preferences for redistribution as well as other topics.

Our study is also related to other strands in the literature; for example, research on tax compliance (Slemrod, 2018; Holz et al., 2020). We contribute by sharing findings from our study of a form of tax compliance that has been entirely overlooked in the literature: tax appeals. Our study also relates to the literature on behavioral public finance showing that insights from behavioral economics can have important implications for public policy (Chetty, 2015). More precisely, we show that information frictions and social preferences play a significant role in the decision to protest taxes.

The rest of the paper proceeds as follows. Section 2 describes the institutional context. Section 3 discusses the role of private benefits. Section 4 presents evidence on the private costs. Section 5 discusses the results on conditional cooperation. 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, 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 DCAD. For each home in the county, the data include information on ownership, address, property characteristics (e.g., number of bedrooms), and historical yearly data on proposed and certified market values, exemption amounts, taxable values, tax rates as well as details on property tax protest records. Whenever needed, we complement the administrative records with other data sources.

The tax amount due is calculated by DCAD using a multi-step formula that starts with the proposed assessment value of the home. Taxes due are a function of a host of factors such as the household's exemptions and the specific tax rates that pertain to the household, depending on the jurisdictions to which the home belongs.[12] Homeowners have the right

[12] The jurisdictions in Dallas County are four county-level jurisdictions, 23 of the cities within the county, nine school districts, the Dallas Community College system, the Parkland Hospital system, and 49 Public Improvement Districts (PID) that fund special public services in some areas (Source: `https://www.dallascounty.org/departments/tax/jurisdictions.php`). Appendix A.1 provides a additional details about the property taxes in Dallas County, while Appendix A.1 provides further details about the

to protest if they disagree with DCAD's proposed assessment value. Among other reasons, homeowners can protest if they believe the proposed value of their property is too high relative to the market value of comparable houses that were sold in the county; if their properties' proposed values are too high relative to the proposed values of comparable houses in the district; or if there are errors in the public records of the property (e.g., an incorrect number of bedrooms). For instance, according to the 2020 data for households that protested on their own, 91.87% of them selected the option "Value is over market value" in their online forms. When filing a protest, homeowners can also provide an "Opinion of Value", which is how much they believe their property was actually worth as of January 1st.

Homeowners can file protests on their own. We refer to this type of protest as *direct protests*. Because such protests are the focus of this paper, we always refer to this type of protests unless we explicitly note otherwise. Instead of protesting on their own, homeowners can hire an agent to protest on their behalf. In exchange for representation, agents normally charge some combination of a flat fee and a percentage of the tax savings (which can be as high as 50% of the tax savings).

For the sake of completeness, we present results for both direct protests and protests through agents; however, protests through agents are less relevant to our study for a couple of reasons. First, we designed parts of the field experiment specifically to reduce the hassle costs from protesting directly – which, if anything, should crowd-out protests through agents. Second, the timing of the protests through agents makes it more difficult for them to be affected by the type of quasi-experimental and experimental variation used in our research design. According to anecdotal accounts, households often sign contracts with agents months before the proposed values are announced. Indeed, the decision to protest through an agent may have been made years previously, as agents offer long-term contracts to automatically protest on the owner's behalf every year.[13]

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 with homeowners the proposed values through its website and, for most households, by mailing a "Notice of Appraised Value". Households have a month from the notification date to file a protest. DCAD's notifications include estimated taxes, which are based on each property's proposed value. These taxes are "estimated" because, technically, property tax rates are determined later in the year, so the county uses the prior year's jurisdictional

data sources.

[13] While there is no 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.

tax rates to estimate taxes due in the Notice of Appraised Value. In practice, tax rate changes are not common so the approximation error is often negligible. For the sake of brevity, in this study we refer to taxes directly, but all analysis is technically based on "estimated" taxes.

In 2020, DCAD presented the proposed values on May 15$^{\text{th}}$; as a result, the deadline to protest was June 15$^{\text{th}}$. After that deadline, disputes can be resolved by homeowners at different stages. Some protests are resolved because the owner accepts the settlements proposed by the county. This settlement may be offered through informal channels, such as an email or phone exchange with a staff member from DCAD. If an agreement is not reached, the protest advances to a formal hearing with a quasi-judicial entity called the Appraisal Review Board.[14] The formal hearing entails no risk: if the DCAD schedules a hearing and households do not attend, the protest is simply dismissed with no penalty.[15] After protests are resolved one way or another, the final assessed home values (from hereon, "certified" values) and tax amounts are calculated, and taxes are due and payable on October 1$^{\text{st}}$, 2020.

In 2020, 8.40% of homeowners protested directly. An additional 8.42% protested through an agent, resulting in total protest rate of 16.82%. This rate of protests is not atypical when compared to recent years: for example, looking at the same sample of households, we find that 13.82% protested in 2017, 15.09% protested in 2018, and 13.89% protested in 2019. Filing a protest directly is simple. Homeowners can protest using a paper form provided by Dallas CAD to households that received notification by mail, a form from the Texas Comptroller that can be printed from the Internet, or via a simple online tool called uFile. To protest online, households need to search for their own name or address on a website, click on their account, and then follow some straightforward steps. In 2020, 75% of direct protests were filed online while the remaining 25% were filed by mail.

Before moving on to the causal identification, we provide some simple descriptive statistics to illustrate the importance of causal identification in this context. Figure 1.a shows the relationship between the tax burden (i.e., the tax amount the household would pay based on the proposed value) and the protest probability. There is a steep relationship: the protest probability is much higher for households that pay higher tax amounts, which holds true for the overall protest rate as well as separately for the direct protest and the protest through

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

[15] Of the 2019 direct protests that contain information on the form in which it was resolved, we find that 31.4% were settled informally, 44.4% were settled after a formal hearing, and 24.2% were either withdrawn or dismissed.

agents. However, this steep relationship could be due to different mechanisms and thus have different implications for our understanding of preferences for redistribution. Specifically, this relationship could be due to selfish reasons: households with a higher tax bill may protest more just because their expected benefits are greater than their expected costs. Alternatively, the steep relationship may be due to social preferences: the richest households protest because they feel that they are contributing more than their fair share. This study explores each of these mechanisms.

Disentangling why some households are more likely to protest can also be of interest to policymakers. For instance, Figure 1.b compares the protest rates of Hispanics and Whites. This analysis is based on a subsample for which we have a proxy for ethnicity based on last names (Perez-Truglia and Cruces, 2017). The results are broken down by each quintile of the distribution of the tax amount in this sample. Note that within each of the first four quintiles is a statistically significant and economically large difference in protest rates between Hispanics and Whites. For example, in the lowest quintile, the probability of protesting is 36% smaller for Hispanics than for Whites (4.99 pp versus 7.84 pp). Across the five quintiles, Hispanic households are 1.97 pp less likely to protest than comparable White households.[16] This statistic constitutes suggestive evidence that the protest system may disadvantage minority groups (Avenancio-León and Howard, 2019). Parsing out the mechanisms behind the decision to protest may serve to explain the roots of these disadvantages. More importantly, the causal estimates may provide hints on how to alleviate the inequalities in the system.

3 Private Benefits

3.1 Conceptual Framework

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

In practice, the amount of taxes that a household pays is calculated through a formula that involves the proposed value and the tax rates for the various jurisdictions as well as other factors, such as the homestead cap and other exemptions. Because households have

[16] This finding is based on a comparison of the 2020 protest rates between Hispanic and White households that control linearly for the logarithm of the 2020 estimated taxes.

the opportunity to protest every year, dynamic considerations may arise, too. For the sake of simplicity and to fix the intuition for the empirical analysis, however, we now introduce a simple model of the decision to protest. Let A be the proposed value of the household and T be the amount the household has to pay in property taxes. Under a simple proportional tax rate (τ), the tax burden without a homestead cap is the following:

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

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

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

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

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

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

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

This equation can be re-arranged as follows:

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

Note that first two terms in the RHS in equation 2 are identical to the first two terms on the RHS in equation 5. Thus, the last term in equation 5 is the difference in incentives to protest introduced by the cap. The cap reduces the expected benefits from protesting when it is binding. Note that the expected benefits are lower the larger the difference between the proposed value and the homestead threshold ($A - \bar{A}$). The intuition is straightforward: absent a cap, the household knows that 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 will 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 universe of 423,607 single-family homes that were subject to property taxes in Dallas County in 2020. In Appendix A.3, we present detailed descriptive statistics for this sample, which we summarize below. The average home has a value of about \$300,000 and pays \$6,150 annually in property taxes, which implies a tax rate of around 2%. For a subsample of these subjects, we obtained individual-level demographic data from a private company.[17] The average subject is 52 years old, 65% are White, 9% are African-American and 20% are Hispanics. Most relevant for this analysis, 74% of these households had a homestead status approved for 2020, and thus their homestead caps can be binding. We focus on this 74% subsample for the main analysis, but leverage the remaining 26% of the households for a falsification test.

Figure 2 summarizes the main results. This figure is a binned scatterplot of the relationship between a given outcome and the distance between the proposed value (A) and the potential homestead cap threshold ($\bar{A}$). For the sake of simplicity, Figure 2 includes a minimal set of controls (the proposed value, a dummy for whether the household protested in the previous year, and a set of school district dummies). This figure also focuses on a narrow band around the homestead cap threshold (\$15,000 above and below). Later we show that the results are robust under alternative specifications. The three panels on the left of Figures 2 (2.a, 2.c and 2.e) correspond to the properties with homestead status – for which the homestead cap threshold can be binding. In turn, the three panels to the right correspond to the properties without homestead status, for which the homestead cap threshold should be irrelevant and thus serve as a falsification test.

We start with Figure 2.a, in which the outcome variable measured on the vertical axis is the amount of property taxes in 2020. The horizontal axis measures the distance to the homestead cap threshold. The blue dots correspond to the observations to the left of the homestead cap threshold, with the blue line corresponding to the linear fit. The coefficient

[17] The company used the names and addresses to merge the records at the individual level. For more details about these data, see Appendix A.2.

from the linear regression is reported in blue, too. For ease of exposition, we normalize all of the coefficients so that they correspond to the effects from a \$10,000 increase in the proposed value. In turn, the red dots and red lines correspond to the observations to the right of the potential homestead cap threshold.

As expected, there is a sharp kink at the threshold: after hitting the homestead cap threshold, households had to pay less in taxes than they would have needed to pay absent the homestead cap. This kink is not only large in magnitude, but also statistically significant: we can reject the null hypothesis that the coefficient to the left of the threshold (-97) is equal to the coefficient to the right of the threshold (-306), with a p-value<0.001. The following thought experiment can illustrate the magnitude of the effect of the cap. Suppose the proposed value starts right at the threshold, and then it increases by \$10,000. The dashed blue line projects the linear fit from the left side to estimate what the outcome would have looked like in the absence of the homestead cap. In contrast, the solid red line shows what the taxes actually looked like under the homestead cap. The estimates indicate that, absent the cap, the \$10,000 increase in proposed value would have resulted in an additional \$209 (=306-97) in taxes. In turn, Figure 2.c reproduces Figure 2.a, but using the tax rate instead of the tax amount as the dependent variable. The results are similar: because of the homestead cap, a household that is \$10,000 above the threshold ends up paying a tax rate that is 0.118 pp (=0.163-0.045) lower.

However, the effects of the homestead cap on taxes (Figures 2.a and 2.c) are mechanical. What we really care about is whether the lower tax burden affects the probability of protesting. To explore that, Figure 2.e is identical to Figures 2.a and 2.c, except that the vertical axis is the protest rate: for example, an indicator that takes the value 100 if the owner protested directly and 0 otherwise. We find that, as expected, there is a sharp kink in Figure 2.e at exactly the homestead cap threshold. This kink is not only large in magnitude, but also statistically significant: we can reject the null hypothesis that the coefficient to the left of the threshold (2.881) is equal to the coefficient to the right of the threshold (-1.604), with a p-value<0.001.

Next, the three panels in the right half of Figure 2 (2.b, 2.d and 2.f) provide a sharp falsification test. These panels correspond to the properties without homestead status, and thus for which the homestead cap threshold should be irrelevant.[18] Thus, this sample provides a natural falsification test: a kink in the right panels would suggest that the results shown above for the left panels are due to a confounding factor. As expected, we find that, in each of the three right-side panels of Figure 2, there are no kinks at the homestead cap threshold:

[18] In these figures, the hypothetical homestead cap threshold is defined as 110% of the assessed value in the previous year (2019).

we cannot reject the null hypotheses that the coefficients are equal below and above the homestead cap threshold. Most importantly, the coefficients are precisely estimated in the right-side panels of Figure 2, meaning that we can rule out not only the large kinks shown in the left-side panel of Figure 2, but also small kinks.

Now, we need a way of quantifying the effects of the homestead cap presented in Figure 2. According to the model from Section 3.1, above, each additional dollar in which the tax bill is reduced due to the homestead cap reduces the expected benefits from protesting, which should translate into a lower protest probability. Indeed, this prediction is borne out by the data: as we move to the right of the threshold, the tax rate decreases linearly (Figure 2.c). Likewise, the protest probability goes down linearly (Figure 2.e). Combining these two figures, we can quantify the relationship between the tax rate and the protest probability. Figure 2.c indicates that going \$10,000 above the homestead cap causes, on average, a reduction of 0.118 pp in the tax rate. In turn, Figure 2.e indicates that going \$10,000 above the homestead cap causes a reduction in the protest probability of 4.485 percentage points. By taking the ratio of these two estimates, we conclude that each 0.1 pp reduction in the tax rate decreases the protest probability by 3.65 pp. To simplify the interpretation of this magnitude, we can express it as an elasticity. The 0.1 pp reduction in the tax rate corresponds to 5% of the average tax rate of approximately 2%, while the 3.65 pp reduction in the protest probability corresponds to 36.7% of the average protest probability. These percent changes suggest that there is a 7.34 ($= \frac{36.7}{5}$) elasticity of the protest probability with respect to the tax rate. In other words, the decision to protest is highly elastic to the financial benefits.

We can translate the previous analysis into a single parameter from an instrumental variable regression. Consider the sample of subjects with homestead status. Let P_i^{2020} be an indicator variable that takes the value 100 if household i protested directly in 2020.[19] We continue the notation from Section 3.1 and define A_i as property i's proposed value and $\bar{A}_i$ as the potential homestead cap threshold (i.e., 110% of the appraised value from the previous year). The relevant Instrumental Variables regression is:

$$\begin{aligned} P_i^{2020} =& \gamma_0 + \gamma_1 \cdot \tau_i + \gamma_2 \cdot (A_i - \bar{A}_i) + X_i \gamma_X + \xi_i \\ \tau_i =& \delta_0 + \delta_1 \cdot \mathbb{1}(A_i > \bar{A}_i) \cdot (A_i - \bar{A}_i) + \delta_2 \cdot (A_i - \bar{A}_i) + X_i \delta_X + \chi_i \end{aligned} \tag{6}$$

X_i stands for an extensive set of additional control variables to improve power and to address any remaining concerns about confounding factors. 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

[19] In all the analysis presented in this paper, we include protests that were marked as received by DCAD through July $15^{15\text{th}}$. For more details, see Appendix A.2.

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.[20] Note that the endogenous variable in the equation (6) is the tax rate, and the excluded instrument is the interaction $\mathbb{1}(A_i > \bar{A}_i) \cdot (A_i - \bar{A}_i)$. The coefficient γ_1 corresponds to the key quantitative exercise from the graphical analysis above: it measures the effect of an additional 1 pp reduction in τ_i, through the cap, on the probability of protesting.

The results are reported in Table 1. Column (1) indicates that a 0.1 pp increase in the tax rate causes an increase of 3.65 pp in protest probability. This coefficient is not only large but also highly statistically significant (p-value<0.001). Figure 2 offers a first falsification test, by reproducing the analysis for properties that are not subject to the homestead cap. Column (2) of Table 1 offers an alternative falsification test: the regression is identical to that from column (1); but the dependent variable indicates whether the household protested in 2019 (instead of whether the household protested in 2020). This is a falsification test in the spirit of event-study analysis: whether the 2020 proposed value ends up being above or below the 2020 homestead cap should not affect whether a household protested a year prior, in 2019. As expected, we find that the coefficient is close to zero (-5.740), statistically insignificant, and precisely estimated (indeed, it is more precisely estimated than the corresponding coefficient from column (1)). We can also confidently reject that the coefficient from column (2), -5.740, is equal to the coefficient from column (1), 36.507, with a p-value<0.001. Moreover, we find that this falsification test is robust not only when looking at the 2019 protests but also when using the protests for all the years for which we have data: 2018, 2017, 2016, and 2015 (results reported in Appendix A.4).

As shown in Figure 2.a, the homestead cap affects the tax rate as well as the tax amount. As a result, for a robustness exercise we can reproduce the analysis using the tax amount instead of the tax rate as the endogenous variable.[21] Columns (3) and (4) reproduce columns (1) and (2), but using the tax amount instead of the tax rate as endogenous variables. The results are robust both qualitatively (i.e., in terms of sign and statistical significance) as well as quantitatively. For example, we showed above that the results from column (1) imply an elasticity of 7.34. If we reproduce that analysis but based on the estimates reported in column (3) instead, we find an elasticity of 13.01, which is somewhat larger but in the same order of magnitude.[22] In turn, column (4) reproduces the falsification test that uses the

[20] Note that this specific regression has no variation in homestead status so that the control variable is irrelevant, but the homestead status variable will have variation in other regressions below.

[21] In principle, households may care about the tax rate, the tax amount, or a combination of both. Whether one or the other is more important in practice may depend on several factors; for example, the tax amount could be more relevant if the protest costs are fixed with respect to the proposed value.

[22] A \$100 increase in the tax amount, which corresponds to a 1.69% increase with respect to the average tax

protest in 2019 as a placebo outcome: as expected, the coefficient is close to zero, statistically significant, and precisely estimated.

In addition to the direct protests, we can look at the effects of the homestead cap on the protests through agents. As discussed in Section 2, above, the timing of the protest through agents is quite different. It is possible that those who protest through agents have signed their contract way before May 15, when the proposed values were announced. As a result, it would be impossible for the homestead cap to affect the decision of those households. It is still possible, however, that the homestead cap could dissuade some households from finding an agent at the last moment. The results are presented in column (5), which is identical to column (1), except that the dependent variable indicates whether the household protested through an agent (instead of indicating whether the household protested on its own). The effects on agent protests are qualitatively consistent with the results on direct protests; however, consistent with the timing issues described above, the effects are quantitatively smaller for the protest through agents. For example, while a 0.1 pp increase in the tax rate would increase the probability of protesting directly by 3.65 pp, a 0.1 pp increase in the tax rate would increase the probability of protesting through an agent by just 0.61 pp (and their difference is statistically significant, p-value<0.001). In turn, the dependent variable from column (6) indicates if the household protested at all, regardless of whether it was direct or through an agent. The effects on total protests (column (6)) are similar in magnitude to the effects on direct protests (column (1)).

Columns (7) and (8) reproduce column (1) for alternative bandwidths. Note that the results should not be identical in magnitude: since the instrumental variables model estimates a local average treatment effect, if there are heterogeneous effects then it would be natural for the estimates to be quantitatively different in different samples. However, we would expect the results to be qualitatively robust and remain in the same order of magnitude. The specification from column (7) is identical to that of column (1) except that it uses a wider band: proposed values must be within \$30,000 of the homestead cap threshold instead of within \$15,000. As a result, the total number of observations increases from 96,274 in column (1) to 179,452 in column (7). The results are both qualitatively and quantitatively quite similar: the coefficient on τ_i is 36.507 (p-value<0.001) in column (1) versus 30.644 (p-value<0.001) in column (7). In column (8), we use an even wider bandwidth: observations within \$150,000 of the threshold, which further increases the sample size to 308,000 households. The results are quantitatively somewhat smaller, but still highly statistically significant (p-value<0.001) and, due to the larger sample size, the coefficient is more precisely estimated.

amount (\$5,916), results in an increase in the protest probability of 2.2 pp, which is in turn equivalent to 22% of the average protest probability (9.98 pp).

4 Private Costs

4.1 Conceptual Framework

If households are able to free-ride by protesting taxes, why doesn't everybody do it? As illustrated in the model from Section 3, one explanation is that there is a private cost associated with protesting that–if significant enough–may discourage households from protesting even if they would like to pay less in taxes. In this section, we seek to provide direct evidence of those costs as well as to quantify them.

There are no fees or monetary costs associated with filing a protest directly in Texas. Instead, the private cost that we are referring to is non-pecuniary. We hypothesized two specific sources of hassle costs, and then designed a mailing intervention aimed at reducing each of those two types of costs. If those costs are significant, our intervention should increase the probability that the subjects protest.

The first source of hassle cost consists of identifying the steps to complete the protest process. While filing a protest on your own is smooth in theory, it may be otherwise in practice. Some households may not even know where to start. Some may think the process is a lot more difficult than it actually is. Other households may be less sophisticated, and thus need step-by-step guidance on how to protest. Indeed, instructions on how to protest are not readily available. At the time of the experiment, only one official source online had instructions on how to file a protest: a PDF document posted on the Dallas CAD's website.[23] However, this document was long, had broad instructions, and was tucked deep into the Dallas CAD's website. There were also a few unofficial online sources, such as blog posts, but those were usually incomplete, outdated, and difficult to find. Moreover, those sources often had a commercial interest, deliberately presenting the protest process more complicated than it really was.

The second source of hassle cost we identified consisted of a specific step in the protest procedure: providing an opinion of value for the home and an argument supporting it. While this process may be undertaken in different ways, typically protesters identify a comparison property that has been sold recently for less than the proposed value of their own property. The comparable property's sale price can then serve as the opinion of value, and information about the recent transaction can be used as the argument. To find a proper comparison property entails a number of steps: First, the household needs to access a tool, such as Zillow, to identify properties that have been sold recently. Then, the household needs to use the tool to filter among recently sold properties those with comparable features yet were sold for less

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

than one's own proposed value within a few months of the start of the year. Though finding this comparison property manually may be relatively easy for someone skilled in searching for information online, this task could be daunting for some people who have limited Internet skills or financial literacy, or possess those skills but not the patience to do the research, or are unfamiliar with the tool. 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, 2020).

It must be noted there are probably other sources of hassle costs in addition to the two sources described above. Even with the aid of our mailing intervention, households must spend time filing and keeping tabs on their protests. Additionally, some may find paperwork to be such an unpleasant activity that its costs well surpass the opportunity cost of time. In that sense, our estimates will provide a lower bound of the full private costs from protesting.

4.2 Mailing Design

We begin with a general description of the letters sent to the subjects. We included a number of features to signal that, though this letter was unsolicited, it came from a legitimate source. The envelope (see Appendix B for a sample) includes the logo of the University of Texas at Dallas, a well-known institution in Dallas County, and the name of one of the professors from that university. The envelope also included non-profit organization postage.

Figures 3 and 4 show the first and second page of this sample letter, but with the addition of some red boxes highlighting the parts that were randomized.[24] The letter contained additional measures to reassure the recipient that the communication was legitimate: the official logo of the University of Texas at Dallas in the header as well as a physical address that they could write to and the URL for the study's website. The website included general information about the study (without discussing any hypotheses or what the study was about) as well as contacts for the researchers and the Institutional Review Board. The letters were tailored to the recipients too: the salutation at the top of the first page included the name of each recipient; and their names and addresses were shown at the bottom of the second page (which appeared through the envelope window).[25]

There are some additional features of the letter that we summarize briefly for now but will be discussed in detail in Section 5. All letters included a table in the middle of the first page with information related to the subjects' properties' proposed values and estimated taxes; we defer the discussion of this treatment arm to Section 5 below. At the bottom of the first

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

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

page, all letters included a URL to an online survey. To verify that the respondents were legitimate subjects and to link survey responses at the household level, we included a unique five-letter survey code for survey access. From here on, we refer to this as the *Field Survey*. The first goal of this survey was to provide a proxy for the dates that recipients opened the letters (Perez-Truglia and Cruces, 2017; Bottan and Perez-Truglia, 2020). The survey included some questions meant to be used as outcomes in the analysis and are discussed in Section 5, below.

4.3 Experimental Design

Subjects can be randomly assigned to receive no letter or to receive one of two types of letters. The *basic aid letter* was designed to address the first type of hassle cost described in Section 4.1, above. This baseline aid consisted of a number of useful tips to help the recipient file a protest. All of that information is found on the *first* page of the letter, a sample of which appears in Figure 3. A key part of the first page is that it mentions recipients could find instructions on how to file a protest on the project's website. We designed our website instructions to be concise, easy to follow, and as explicit as possible. A full copy of the website is in Appendix D. The website included step-by-step instructions on how to file a protest online or by mail. These walkthroughs included hyperlinks to the relevant websites as well as screenshots of a fictitious application for added clarity.

The second letter type, *extra aid letter*, is identical to the basic aid letter except that it includes additional information on the second page. Figure 4.a shows what the second page of the letter would look like if assigned to the basic aid treatment, while Figure 4.b shows what it would look like under the extra aid treatment. The extra aid message in Figure 4.b is highlighted inside of a red box with dashed lines (this box is shown for expositional purposes and was not included in the actual letters sent to subjects). The extra aid message was intended to reduce the second source of hassle costs described in Section 4.1, above.

The first paragraph of the extra aid message began by providing some facts about the protest filing process, such as that protesting is simple, can be done without an agent, may not require a hearing (which could be intimidating to some subjects) when the DCAD proposes a settlement offer; and even if a hearing is scheduled, there is not any risk if it is not attended. Then, it provided an argument for the protest. More specifically, we presented the most common type of argument: based on the recent sales price of a comparable property, the proposed value for the property is over the market value.[26] To make it simpler to use this information, we presented it much how it would look on the actual protest form: with

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

a check mark in the "Value is over market value" box, an opinion of value field with the sales price of the comparison property filled in, and then a handwritten note with a usable argument. For example, in the sample letter shown in Figure 4.b, the handwritten note reads, "I found a home that is similar to mine but was recently sold for less than my home's appraised market value. The property located at 2234 Meadowstone Dr. (Carrollton, TX) is 0.20 miles away from my home, and has the same number of bedrooms and a similar square footage. That property was sold on 10/31/2019 for $160,000." Households could have used our proposed argument directly; however, to clarify that the content is just a suggestion, we included the following: "You can find information about this sale by searching for the property's address on Zillow.com or Redfin.com. On these websites you can find other comparable properties to support your protest." Additionally, we mentioned that subjects could protest based on different arguments, offering the following message: "You can also 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 websites such as Zillow and Redfin. For each subject, the algorithm searches for properties that have been sold by late 2019 or early 2020, and were similar to the subject's own property in a number of dimensions (e.g., number of bedrooms, bathrooms, square footage built, location) but were sold for less (between 5% and 20%) than the proposed value for the subject's own property. In Appendix A.5, we provide details about this algorithm as well as some descriptive statistics. If a homeowner were to hire an agent to protest on his or her behalf, we believe it is likely that the agent would use a similar (or even the exact same) argument.

4.4 Subject Pool and Implementation Details

We started with the universe of 423,607 residential single-family properties used for the analysis in Section 3.2, above, and focused on a subsample of 78,462 of those households, which constitutes our subject pool for the field experiment. We arrived at that subsample after applying a number of filters. For example, we excluded households that, according to the latest available data from DCAD, had already filed a protest – because it would not matter if we sent them a letter or what we included in the letter, because their decision had been made already.[27] The most important condition was to focus on households for whom our algorithm could find comparison properties that we could utilize in the extra aid message –

[27] We initially selected a sample of 79,322 properties. However, due to a lag of a few days in the way DCAD reports the data, we had to drop 860 of them from the subject pool; with the updated data, we discovered that they had already protested by the time we mailed the letters.

for the full criteria, see Appendix A.2. In Appendix A.3, we provide the descriptive statistics for the subject pool. While not identical, the subject pool is quite similar to the universe of households in observable, pre-treatment characteristics. Additionally, Appendix A.7 breaks down the pre-treatment characteristics by treatment groups. We find that, consistent with successful random assignment, the observable characteristics are balanced across all treatment groups.

The timing of the intervention was carefully planned.[28] We created the letters as soon as the administrative data including the 2020 proposed values became available (on May 16$^{\text{th}}$, 2020). To accelerate delivery time, we used a mailing company located in Dallas County (i.e., within the same county as all of the recipients). As a result, the vast majority of the letters should have gotten to the subjects quickly. The mailing company dropped the letters off at the local post office's facility on May 20$^{\text{th}}$. This company estimated that the vast majority of the letters would be delivered in the next couple of days. Consistent with this projection, we began to receive some answers to the Field Survey and visits to the website on May 21$^{\text{th}}$. Moreover, the post office scans mail pieces when they reach the last mile immediately before delivery. More than 90% of the letters had been scanned by Friday, May 22$^{\text{nd}}$, 2020. Based on data from previous years, most subjects file protests close to the deadline, which in 2020 was June 15$^{\text{th}}$. For that reason, we feel confident that there was enough time between receipt of the letter and the protest deadline so that, at least for most households, the information provided in the letter could influence their decision to protest. Moreover, in Appendix A.8, we provide suggestive evidence in support of this view based on the dates when subjects responded to the Field Survey and visited the project's website.

4.5 Econometric Model

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

$$P_i^{2020} = \eta_0 + \eta_{basic} \cdot L_i^{basic} + \eta_{extra} \cdot L_i^{extra} + X_i^{pre} \eta_X + \epsilon_i \tag{7}$$

The variable, L_i^{basic} is an indicator that takes the value 1 if the household was mailed a basic aid letter and 0 otherwise. Likewise, L_i^{extra} is an indicator that takes the value 1 if the household was mailed an extra aid letter. Lastly, X_i^{pre} is a vector of pre-treatment controls,

[28] For additional details about the implementation of the field experiment, see Appendix A.6.

which is the exact same set of control variables listed in Section 3.2, above. 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, in the spirit of event-study analyses, we use the pre-treatment data to construct falsification tests by using protest in previous years as the dependent variable.

4.6 Results

The regression results are presented in Table 2. All regressions are based on the same specification given in equation (7), above, but differ according to the dependent variable. The dependent variable in column (1) takes the value 100 if the owner protested directly and 0 otherwise. The basic aid letter increased that probability of protesting by 1.792 pp, a highly statistically significant (p-value<0.001). In turn, the extra aid letter increased the protest probability even more, by 3.509 pp.

Column (2) of Table 2 presents an event-study falsification test. This column is identical to column (1) except that the dependent variable indicates whether the owner protested directly in 2019 instead of 2020. Since the letters were sent in 2020, they should have no effect on the 2019 protests. As expected, we find point estimates that are close to zero, statistically insignificant, and precisely estimated. Moreover, we find that this falsification test is robust not only when looking at the 2019 protests but also when looking at protests in all the years for which we have data: 2018, 2017, 2016, and 2015 (results reported in Appendix A.4).

Column (3) shows the effects on the probability of protesting through an agent. Since our letter is providing aid for households to protest directly, it should not make them more likely to protest through an agent. If anything, our letter could crowd-out protests through an agent: for example, households that were about to hire an agent may receive our letter and decide to protest directly instead. However, the data suggest that there was no such crowd-out: the coefficients on the letters are close to zero (0.030 and -0.122 for the letters with the basic aid letter and extra aid letter, respectively), statistically insignificant, and precisely estimated. A likely explanation for why there was no crowd-out is based on the timing. As explained in Section 2, above, it is likely that households had signed contracts with agents earlier in the year or even long-term contracts in previous years. As a result, by the time they received our letter, it would have been too late for them to change their decision. In turn, the results in column (4) combine both types of protests: the dependent variable in column (4) indicates whether the owner protested, either directly or through an agent. Due to the lack of effects on agent protests (column (3)), the effects on total protests (column (4)) are almost identical to the effects on direct protests (column (1)).[29]

What were the precise mechanisms behind these effects? Column (1) indicates that the effect of the basic aid letter was positive (1.792 pp) and statistically different from the effect of the extra aid letter (3.509 pp). The difference between the coefficient estimates indicates that the extra aid message, on its own, had an effect of 1.717 pp (= 3.509 – 1.792), which was highly statistically significant (p-value<0.001). The content of the aid message was quite specific, so the interpretation of its effect is straightforward. In contrast, the basic aid letter had many components to it, so it is a bit less clear what the underlying mechanisms were.

One potential interpretation is that the basic aid letter acted as a reminder of the opportunity to protest, or made it more salient. This explanation is unlikely, however, as the proposed property taxes are quite salient around the time subjects received our letter. To test this hypothesis more directly, we exploit heterogeneity on whether households were mailed a notification by the DCAD. Starting on May 15, any homeowner was able to download the Notice of Appraised Value by going to the DCAD webpage (we provide a sample of this notification in Appendix G). Additionally, DCAD sent notifications by mail to some households but not to others. For example, all households whose proposed value increased relative to the previous year were mailed a notification. DCAD mailed the official notifications on Friday, May 15, a few days earlier than we mailed our letters, on Wednesday, May 20, so the households should have received the official notification around five days before our letter. If our basic aid letter worked primarily through a reminder effect, it should have had a larger effect on the households that did not receive the Dallas CAD notification.

The results are presented in columns (5) and (6), which split the sample based on whether the subjects were (column (5)) or were not (column (6)) mailed a notification from the DCAD.[30] The effects of the basic aid letter on subjects that were and were not mailed a notice (coefficients of 1.449 and 1.935, respectively) are statistically indistinguishable from each other. If anything, the point estimate is a bit larger in the sample that received a notification from the DCAD – that is, the difference has the opposite sign as the one predicted by the reminder channel. This finding constitutes suggestive evidence that the effect of the basic aid letter went way beyond a simple reminder effect. As additional evidence, in Appendix A.9, we present evidence from a Regression Discontinuity Design showing that the official DCAD mail notification did not have significant effects on the probability of protesting. If the official notification did not act as a reminder, it is unlikely that our unofficial notification would.

Instead, our favorite interpretation is that the walkthroughs that we provided through our

[29] The only difference between the results in columns (1) and (4) lies in the effects relative to the baseline (i.e., the ratio between the coefficient and the corresponding average of the dependent variable).

[30] We split the sample using publicly available data that DCAD posted on its website; for details, see Appendix A.9.

project's website were one of the key mechanisms behind the effects of the basic aid letter. A first piece of evidence for this interpretation relies on a perhaps unusual source of data. In the project's website, we provided an email address to contact the researchers in case the subjects had any concerns about the research project. While this was not the purpose of providing the email address, a number of subjects sent emails to this address expressing gratitude for the letter. We cannot share these messages because they are confidential, but their content is broadly consistent with the letter having reduced subjects' hassle costs. For example, some subjects mentioned that they had wanted to protest for years but did not know how to until receiving our letter. Other subjects mentioned that they thought protesting was more complicated than it was until they looked at our instructions. Similarly, the Field Survey included a final, open-ended question in case the subjects wanted to share any thoughts with the researchers. Many subjects used that space to express gratitude and sometimes made explicit how the information contained in the letter and the website helped them navigate the protest process.

A second piece of evidence for this mechanism relies on data from the traffic to the project's website. We start by noting that the basic aid provided in the letters generated a total of 903 additional direct protests; 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). We can compare this number of additional protests to the number of unique visits to the website. Google Analytics recorded a total of 2,769 unique visits to the walkthroughs to protest online or by mail (for more details, see Appendix A.8). Some of those visitors may have looked at the walkthroughs but did not protest, while some of those visitors may have used the walkthroughs but would have protested even without them. If we assume that around a third of those website visitors were induced to protest by our website, we would explain all the 903 additional protests generated by the basic aid information. In other words, it would not be far-fetched to attribute the entire effect of the basic aid letter to the walkthroughs.

For more evidence on how our website played an important role in the effects of the letters, notice that, of the households that visited the walkthroughs, 94.8% looked at the online walkthrough and the rest looked at the mail walkthrough. As a result, we would expect that the effects of our letters acted primarily through the online protests. To test this hypothesis, in column (7), the dependent variable is an indicator that takes the value 100 if the household protested directly online but 0 if the household protested by mail or did not protest at all. The evidence suggests that the vast majority of the effects operated through online protests: that is, the coefficient on the basic aid letter from column (7), 1.591, is 88.7% as large as the corresponding coefficient from column (1), 3.509. In other words, and consistent with the data on the visits to the website, we find that a strong majority (88.7%) of the additional protests induced by our letter were conducted online.

We can also provide some complementary evidence about the effects of the extra aid message. One simple interpretation is that subjects just took our suggested argument and used it "as is" in their protest form. An alternative interpretation is that they used the message as a proof of concept but then figured out an argument on their own. We can conduct a suggestive test of these alternative hypotheses using data from the online protests, for which we can observe the opinion of value that the protesters entered in the online form. The dependent variable in column (8) takes the value 100 if the households provided an opinion of value in their protest, which is within half a percentage point of the value we selected for their extra aid message. However, these results must be taken with a grain of salt: they are based on a non-random subsample (households that protested online and entered a value in the opinion of value field), and thus they may be subject to endogeneity biases despite the random assignment.

Column (8) shows that, in the control group with no letter, there is a 3.37 pp chance that a household enters an opinion of value that almost coincides with the value that we would have shown them if they had been assigned to the extra aid message treatment. In other words, it is a rare coincidence for subjects to use an almost identical opinion of value as the one we would have suggested. For households that received the basic aid letter, that probability remains equally low. For households that received the letter with the extra aid, however, this outcome jumps drastically by 15.287 pp (p-value$<$0.001), which constitutes strong evidence that a substantial fraction of households that received the extra aid message used it "as-is".

Lastly, we want to see if our treatments were consequential. One potential concern is that our letters induced protests that were ultimately not successful – that is, that they were just a waste of time. We can provide direct evidence that this was not the case. In column (9), the dependent variable takes the value 100 if the household protested directly and won (i.e., received a discount in their assessed value) and 0 otherwise (i.e., if the protest was unsuccessful, or if the household did not protest directly). The coefficients are still economically and statistically significant. The ratio between the coefficients from columns (1) and (9) suggests that 52% (55%) of the marginal protests that were induced by our letter with basic (extra) aid were successful. These success rates are comparable to the corresponding success rate of 58% in the control group.[31] In other words, the additional protests that were induced by our letters were, on average, roughly as successful as the protests in the control

[31] This success rate is based on the ratio of the share of direct protests that were successful (5.07, from column (9)) to the share of direct protests (8.67, from column (1)).

group.

In the same vein, column (10) is similar to column (9), except that instead of measuring the extensive margin of the success, it measures the intensive margin. In column (10), we define the dependent variable as equal to the percentage point reduction in the assessed value, regardless of the type of protest. This outcome takes the value 0 by construction for subjects with no protests or with unsuccessful protests. Consistent with the extensive margin results from column (9), we find significant effect on this outcome, too, further corroborating that our letters were consequential.

4.7 Magnitude of the Hassle Costs

In this section, we discuss the economic magnitude of the hassle costs in more depth. A first challenge of interpreting the magnitudes in mailing experiments is non-compliance: for example, some households may not have received the letters, or they may have received them but did not read them. To correct for these types of non-compliance, we need an estimate of the reading rate (i.e., the share of recipients that actually read the letter on time). Following Bottan and Perez-Truglia (2020), we combine estimates from different sources to approximate the reading rate. According to the US Monitor Non-Profit Standard Mail Delivery Study, around 95% of standard non-profit mailers were successfully delivered (U.S. Monitor, 2014). Based on data from the US 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.[32] If we combine the two estimates above, we arrive at a reading rate of 70.3% ($= 0.95 \cdot 0.74$). As a result, to account for this source of attenuation bias, we need to scale the coefficients up by a factor of 1.42 ($= \frac{1}{0.703}$). The resulting scaled-up effects would be 2.55 pp for the basic aid letter and 4.98 pp for the extra aid letter. This is still a conservative scale-up factor, as some households may not have opened the letter on time (i.e., after they had filed a protest or after the protest deadline, whichever came first), but we are not accounting for that form of non-compliance.

To translate the hassle costs into dollar amounts, we combine the results from the field experiment with the results on the homestead cap from Section 3. We focus on the effect of the most complete letter (the extra aid letter), which would still give a lower bound on the costs from protesting, as this letter did not eliminate the hassle costs completely. For example, subjects still had to follow the instructions to file the form, which may take at least a few minutes. Subjects may also need to take further action in the future, such as discussing

[32] 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 (2020) for more details.

a settlement in formal or informal hearings. The scaled-up effect of the extra aid letter is 4.98 pp. According to the results from column (3) of Table 1, we would need to reduce the tax amount by \$226 ($= \frac{4.98}{0.022}$) in order to generate an equivalent reduction of 4.98 pp in the protest rate. That is, our lower bound for the hassle costs are estimated at \$226.[33]

As a sanity check, we can compare our estimate of hassle costs to the fees charged by agents that protest on households' behalf. We identified one such company that offered the service for a flat fee. Let's assume that the marginal client of this firm is indifferent between hiring this agent or protesting on her/his own. In that case, the flat fee should constitute a measure of his/her hassle cost.[34] In 2020, the flat fee was \$139 for properties assessed below \$200,000, and \$305 for properties assessed between \$200,000 and \$500,000. [35] Those flat fees (\$139 and \$305) are in the same order of magnitude as our estimated average hassle cost (\$226), thus suggesting that our estimates are in the right order of magnitude.

5 Conditional Cooperation

5.1 Conceptual Framework

The above sections provide evidence suggesting that households are quite elastic to the private benefits and costs of protesting. While this evidence suggests that selfish motives are important, we show that social preferences ameliorate incentives to act selfishly. More precisely, we measure conditional cooperation: that is, whether households are more willing to tolerate taxes when they think their fellow citizens are also paying taxes. Let τ represent the household's own tax rate, while $\hat{\bar{\tau}}$ represents the household's perception about the average tax rate in the county. We can summarize the decision to protest with the following equation:

$$\mathbb{P}(P = 1) = \alpha_{own} \cdot \tau + \alpha_{avg} \cdot \hat{\bar{\tau}} + \zeta \tag{8}$$

The coefficient $\alpha_{own} < 0$ would represent the effects of the private benefits: for example, households dislike paying taxes, so they are more likely to protest when they face higher tax rates. Most importantly, an $\alpha_{avg} > 0$ would indicate that households are conditional cooperators, that is, they are more willing to tolerate taxes if other households contribute

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

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

[35] Source: `https://www.dfwtaxadvisor.com/practice_areas/property-tax-protest/`, accessed on May 15$^{\text{th}}$ 2020.

too.[36] Moreover, the ratio between the coefficients α_{avg} and α_{own} can quantify the strength of conditional cooperation: that is, what additional tax rate households tolerate for each percentage point increase in what they think everyone else is paying.

5.2 Design of the Field Experiment

We designed a treatment arm of the field experiment to measure the strength of conditional cooperation in the context of property taxes. Ideally, we would just randomize how much others are paying in taxes and measure the effect of that amount on protest rates. However, that measure is not feasible. Instead of manipulating the average tax rates directly, we manipulate the subject's perception of the average tax rate. We use an information-provision design to create exogenous variation in those perceptions. The idea is to leverage that the information on how much others pay in taxes is not easily accessible and thus probably not well known; for example, some households may underestimate how large $\overline{\tau}$ is, while others may over-estimate it. Hence, by providing households with accurate information about $\overline{\tau}$, we can induce exogenous shocks to households' perceptions of $\overline{\tau}$. Then, we can measure whether these information shocks affect their decisions on whether to protest.

We introduced another layer of randomization in the mailing experiment: the content of the table appearing in the middle of the first page. In Figure 3, this table is highlighted inside a red box with dashed lines. This box is just for explanatory purposes (it was not included in the actual letters sent to the subjects). All letters included a table, but we randomized (with 50% probability) whether the table included the column with the figures for the average Dallas home (positioned in the last column). This table provided information on whether households' tax rates were above or below the average Dallas home, and by how much.[37]

The random inclusion of an additional column was meant to provide a shock to households' perception of the tax rate other households pay. However, a naïve comparison of the protest rate among households that received the information shock and those that did not, would deliver incorrect estimates to the extent that households update heterogeneously to the information shock (Cruces et al., 2013; Cullen and Perez-Truglia, 2018; Bottan and Perez-Truglia, 2020; Fuster et al., 2018). For example, suppose that half of the sample was originally downward-biased and, thus, updated upwards, while the other half, who was orig-

[36] The parameter ζ agglomerates the remaining factors, such as the hassle costs. Also, note that this simple model makes the implicit assumption that households care about the tax rate that others pay. In theory, households could care about the average tax *amount* instead of the average tax *rate*. The main motivation for this specification choice is survey data indicating that the vast majority of households think that it would be fair for all households to pay the same tax rate. The results are presented in Appendix A.10.

[37] As described in the RCT pre-registration, we cross-randomized a minor aspect of the table for a robustness check: whether it included a row making the tax rates explicit. For more details about this additional randomization and the results, see Appendix A.7.

inally upward-biased, reacted in the opposite direction. In such a case, despite the fact that households react to the information, we would find a null average effect of information disclosure, because the two opposite effects cancel out. To overcome this limitation, and as anticipated in the pre-registration, we adapt the research design from Bottan and Perez-Truglia (2020) to our context. We explain this in detail below.

5.3 Design of the Mturk Survey

We designed a supplemental survey to validate the research design from the field experiment. This survey was conducted on a separate sample recruited from Amazon's Mechanical Turk (Mturk) online marketplace. From here on, we refer to this survey as the Mturk Survey. The goal of the survey was to reproduce the same information-provision experiment from the field experiment and measure how households updated beliefs in response to the information given to them. One key advantage of the Mturk Survey is that it can measure perceptions both before and after the information-provision takes place, thus allowing us to observe the prior misperceptions as well as the belief updating.[38]

The full survey instrument was included in the same RCT pre-registration, and the survey was conducted around the same dates as the main field experiment: from June 5$^{\text{th}}$ to June 15$^{\text{th}}$, 2020. The full survey instrument of the Mturk Survey is in Appendix F, and is summarized below:

- **Step 1 (Elicit Prior Belief):** We elicit perceptions about average property taxes and average home market values in their county in 2019.

- **Step 2 (Information-Provision Experiment):** All respondents were told that some survey participants will be randomly chosen to receive information about home values and property taxes in 2019. In the following screen, respondents discovered whether they were assigned to receive the information. Just like in the field experiment, subjects were assigned to receive the information with 50% probability.

- **Step 3 (Elicit Posterior Belief):** We re-elicited their guesses about average property taxes and average home market values in their county. Following (Cavallo et al., 2017), we elicited perceptions about average property taxes and market values in their county in 2020 instead of 2019. We did this to avoid asking the same question twice, which could induce respondents to just repeat their initial guess. By eliciting the posterior

[38] Relative to the Field Survey, the Mturk Survey has other advantages, such as the fact that the information-provision experiment happens after the start of the survey, and thus we do not need to worry about endogenous selection into the survey.

about 2020, we also avoid pressuring the respondents to answer exactly with the feedback given to them in Step 2. Moreover, to try to mimic the field experiment as closely as possible, we took the respondents' posterior beliefs and created a table identical to the one used in the field experiment: comparing their own property taxes and homes' market values to those of their respective county averages.[39]

- **Step 4 (Elicit Survey Outcomes):** We measured a number of outcomes that we conjectured could be affected by the perceptions of the average taxes paid by others. The first question simply asked whether the household pays too much in taxes relative to other households in the county. The second question asked whether the household pays too much in taxes without making any reference to the taxes paid by others. The next question elicits this last question in a more quantitative way, by asking what would be a fair amount for this household to pay in property taxes, explicitly holding constant how much other households pay.

We followed several best practices for recruiting individuals in Mturk – for more details about the design and implementation of this survey, see Appendix A.11. We collected responses from 2,065 U.S. homeowners. Appendix A.3 provides descriptive statistics for this sample. We show that, in terms of their observable characteristics, the Mturk sample is certainly not identical to the other samples used in this paper, but not wildly different either.

5.4 Results from the Mturk Survey

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

[39] We also mimicked the additional randomization we included in the field experiment related to whether the table would contain the extra row making the tax rates explicit. For more details, see Appendix A.12.

x-axis in Figure 5.a to update their beliefs downwards, while individuals toward the right side of the x-axis should update their beliefs upwards.

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

Figure 5.b illustrates the intuition behind the identification strategy: we know whether the disclosure of accurate information translated into upwards or downwards revisions in beliefs, based on whether a household is toward the left or the right of the x-axis. We can summarize these findings in a single regression coefficient, based on the adaptation of the disclosure-randomization design of Bottan and Perez-Truglia (2020) to our context. Let τ_i be the individual's own tax rate, $\overline{\tau}$ be the *actual* average tax rate in the county and let $\overline{\tau}_i^{post}$ be the individual's posterior belief about the average tax rate in the county. And let D_i be an indicator variable that takes the value 1 if the information on the average tax rate was shown to the subject. The regression of interest is the following:

$$\overline{\tau}_i^{post} = \nu_0 + \nu_1 \cdot D_i \cdot (\overline{\tau} - \tau_i) + \nu_2 \cdot (\overline{\tau} - \tau_i) + \nu_3 \cdot D_i + \varepsilon_i \tag{9}$$

The coefficient ν_2 shows the relationship between the truth and the posterior beliefs when the truth is not disclosed to the individual, while ν_1 measures how much stronger that relationship becomes when the truth is disclosed. In Figure 5.b, ν_1 would correspond to the difference between the slopes in the treatment versus the control groups. A coefficient $\nu_1 > 0$ would indicate that subjects are incorporating the information shown to them.

Table 3 shows the results from this econometric model based on the data from the Mturk Survey. All regressions in this table are based on the same specification, but using different dependent variables. The dependent variable in column (1) is the posterior belief about the average tax rate in the county. The coefficient on the variable $\overline{\tau}$ shown in the table corresponds to the information shock variable (i.e., $D_i \cdot (\overline{\tau} - \tau_i)$). The coefficient on $\overline{\tau}$ from column (1) indicates that a 1 pp increase in the information shock increases the posterior belief by 0.393 percentage points. This finding demonstrates that subjects update their beliefs when given information. 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 perfectly consistent with rational learning: for example, individuals may not fully update 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. Most importantly, the rate of pass-through is significantly above zero, and in the ballpark of the pass-through found in survey experiments covering a range of topics.[40]

Column (2) provides a falsification test in an event-study fashion. The dependent variable is the prior belief (i.e., measured *before* the information-provision experiment) about the average tax rate in the county. Because the information cannot have effects until it is disclosed, we would not expect an effect on the prior beliefs. As expected, the coefficient from column (2) is close to zero, statistically insignificant, and precisely estimated.

Another potential concern is that the information provision on the average tax rate may affect beliefs other than the average tax rate. In the context of the Mturk experiment, subjects may not remember the exact tax rate that they pay themselves and, as a result, may use the feedback on the average taxes to update their beliefs about their own tax rates.[41] To see if the concern is real, column (3) shows a specification in which the dependent variable is the posterior belief about their own tax rate. The coefficient is close to zero, statistically insignificant, and precisely estimated, indicating that the information shock did not affect the respondents' perceptions about their own tax rates.

The previous evidence validates the design of the information experiment, which was the main goal of the Mturk Survey. Additionally, we can use the Mturk Survey to assess whether the shocks to perceptions about the average tax rate had an impact on attitudes toward taxation. We report these results in columns (4) through (6) of Table 3, each of which uses as the dependent variable a different measure of whether the household thinks that their taxes are too high. In column (4,) the dependent variable reports a *relative* measure: whether the household states that their own taxes are too high specifically in comparison to other households in the county, on a 1–10 scale. In column (5), instead, we report an *absolute* measure: here, the dependent variable takes the value 100 if the respondent thinks that their

[40] For example, Bottan and Perez-Truglia (2020) finds a pass-through of 0.205 in the context of home price expectations (coefficient reported in column (1) of Table 2). And Cavallo et al. (2017) reports that, when forming inflation expectations, the average Argentine respondent assigns a weight of 0.432 to the information and the remaining 0.568 to their prior beliefs (coefficient α-statistics reported in Panel B, column (1) of Table 1).

[41] This concern is limited to the Mturk experiment because all letters we mailed included the tax rate that subjects themselves pay.

own taxes are too high. Finally, in column (6) the dependent variable is the desired tax cut: in other words, the difference between the tax rate that the household currently pays and the tax rate that the household thinks it would be fair to pay. The coefficients from columns (4) through (6) are all negative, indicating that perceiving that others pay more in taxes leads households to complain less about their own taxes. The three coefficients are on the same order of magnitude: a 1 pp information shock decreases the outcome variable by 0.16, 0.10, and 0.24 standard deviations in columns (4), (5) and (6), respectively. Two out of these three coefficients are statistically significant: p-values of 0.060, 0.205, and 0.039, in columns (4), (5), and (6), respectively. In sum, this evidence suggests that households' tolerance for taxation depends on what they think the tax rates paid by others are. In the following section, we measure if these effects are also present in the decision to protest property taxes.

5.5 Results from the Field Experiment

We now turn to the results from the field experiment. This analysis is based on the sample of 50,394 subjects from the field experiment who were randomly selected to receive a letter. For this analysis, we use the same econometric model presented in Section 5.4 above. Recall that P_i^{2020}, the main outcome, is an indicator variable that takes the value 100 if the subject filed a protest in the post-treatment period. The regression of interest is as follows:

$$P_i^{2020} = \nu_0 + \nu_1 \cdot D_i \cdot (\overline{\tau} - \tau_i) + \nu_2 \cdot (\overline{\tau} - \tau_i) + \nu_3 \cdot D_i + \nu_M \cdot M_i + X_i^{pre} \nu_X + \varepsilon_i \tag{10}$$

The key coefficient of interest is ν_1, corresponding to the effect of the information shock. Additionally, to help interpret the magnitude of this coefficient, we include an additional independent variable, M_i, that takes the value 1 if the subject was assigned to receive the extra aid message and 0 if not. Its coefficient, ν_M, measures the effect of the extra aid message. X_i^{pre} corresponds to the vector of control variables, which contains the same variables used for the rest of the analysis and listed in Section 3.2 above. Again, since this is an experiment, the goal of using pre-treatment controls is to gain statistical power by reducing the variance of the error term (McKenzie, 2012). We also use the pre-treatment data to construct falsification tests in an event-study fashion.

Table 4 presents the estimation results. As in Table 3, the coefficient on $\overline{\tau}$ corresponds to the coefficient on the information shock variable (i.e., $D_i \cdot (\overline{\tau} - \tau_i)$). Additionally, to help interpret the magnitude of the coefficient on $\overline{\tau}$, Table 4 reports two additional rows of coefficients related to the findings presented in the previous sections. The second row reports the effect of the extra aid message (i.e., coefficient ν_M. The coefficient on τ_i, reported in the last row of Table 4, corresponds to a separate regression: an Instrumental Variables regression used in the quasi-experiment analysis from Table 1, but restricted to the same sample used to estimate the coefficient on $\overline{\tau}$. Recall that this Instrumental Variable regression exploits variation within the subset of households with homestead status in 2020. And to maximize statistical power, we use a bandwidth of proposed values within \$150,000 of their respective homestead cap thresholds.

The columns of Table 4 correspond to different regressions with the same specification but different dependent variables. In column (1), the dependent variable takes the value 100 whether the household protested directly in 2020. The coefficient on $\overline{\tau}$ indicates that an additional 0.1 pp in information shock about $\overline{\tau}$ decreases the probability of protesting in 2020 by 0.094 pp, which is statistically significant (p-value=0.066). Column (2) presents a falsification test in the spirit of event-study analyses. The dependent variable is the protest indicator for 2019 instead of 2020. Since our letters had not been sent yet, the information shocks should not have had an effect on 2019 protests. As expected, the coefficient on $\overline{\tau}$ from column (2) is statistically insignificant. However, this coefficient is imprecisely estimated, and thus we cannot reject the null hypothesis that it is equal to the coefficient from column (1). We also ran this falsification test using data on protests from 2018, 2017, 2016, and 2015, and found that the estimated coefficient on $\overline{\tau}$ is statistically insignificant for each of those years (results reported in Appendix A.4).

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

Column (3) of Table 4 repeats the same analysis as in column (1), except that it is restricted to the subsample of 1,888 households that responded to the Field Survey.[42] The survey respondents are certainly not a random sample of the subject pool (for a detailed comparison, see Appendix A.3). The most important difference is that the protest probability is much higher in the survey sample (50.26% protested) relative to the full sample (11.29%

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

protested). Our preferred interpretation is that the subjects who paid the most attention to our letter were those on the fence about protesting or not.[43]

When we focus on the survey respondents, the coefficient on $\overline{\tau}$ is -12.566, and statistically significant (p-value=0.021). This coefficient is substantially larger than the coefficient reported in column (1), which is partly mechanical due to the larger baseline rate (50.26 in column (3) versus 11.29 in column (1)). Additionally, we attribute this difference to the fact that the subsample of survey respondents paid closer attention to the letter. Considering that both the *Extra Aid Message* and the average tax rate were included in the letter, we should observe a larger coefficient on *Extra Aid Message* in column (8) as well. Consistent with this interpretation, while the coefficient on $\overline{\tau}$ is substantially larger (13 times as large) in the survey respondents sample (column (3)) relative to the full sample (column (1)), the coefficient on *Extra Aid Message* is also substantially larger (16 times) in column (3) relative to column (1).

One concern with the analysis from column (3) of Table 4 is that the results may be contaminated by endogenous sorting into the survey. To address this concern, column (4) presents the same type of falsification test from column (2), where we used the 2019 protest instead of the 2020 protests as the dependent variable, but restricting the sample to survey respondents. The rationale for this exercise is that observing "effects" on the 2019 protests would suggest that they are driven by selection into the survey. Reassuringly, the effects on that placebo outcome are close to zero, statistically insignificant, and precisely estimated.

The Field Survey included two questions meant to be used as outcome variables: the *unfairness* of the amount of property taxes paid relative to other households in the county, and the stated *intention* to protest property taxes at the time of answering the survey.[44] The effects on those survey outcomes, U_{2020} and I_{2020}, are presented in columns (5) and (6) of Table 4.

In column (5), the dependent variable U_{2020} is whether the households think that their own taxes are unfair relative to the taxes paid by others. We would expect that an increase in the perceived $\overline{\tau}$ would reduce this feeling of unfairness. As expected, the coefficient on $\overline{\tau}$ is negative and statistically significant (p-value=0.060). The magnitude is rather large as well: a 0.1 shock to $\overline{\tau}$ decreases the feeling of unfairness by 0.046 points on a 10-point scale. We can also look at the effects of the *Extra Aid Message* and the homestead cap on this unfairness outcome. Due to its content, we do not expect the *Extra Aid Message* to affect the feeling of unfairness. Indeed, the coefficient on *Extra Aid Message* from column (5) is close to zero (0.083), statistically insignificant, and precisely estimated. Regarding the coefficient on τ_i,

[43] This difference may be driven by reciprocity, too: the households that found our letter helpful were the ones likely to protest and wanted to reciprocate our help by responding to our survey.

[44] A copy of the full survey is included in Appendix E.

we expect it to be positive: the question is whether your own taxes are unfair relative to the taxes paid by others. The homestead cap increases your own tax rate while holding constant the tax rates of everyone else. As a result, the homestead cap should have an effect on this unfairness outcome. As expected, we find the coefficient on τ_i to be positive (1.730) and statistically significant (p-value=0.046).

The dependent variable in column (6), I_{2020}, corresponds to the intention to protest at the time of the survey: it takes the value 100 if the household states that it is likely or very likely that it will protest in 2020, and 0 otherwise. Because the households were less likely to protest, we would anticipate a negative effect on their expected probability of protesting, too. Consistent with this prediction, the coefficient on $\overline{\tau}$ is negative and highly statistically significant (p-value=0.008). By the same token, we would expect the coefficients on *Extra Aid Message* and τ_i to be qualitatively consistent between column (6) (expected behavior) and column (3) (actual behavior). Indeed, we find that those pairs of coefficients are not only qualitatively consistent but also consistent in magnitude.

5.6 Magnitude of Conditional Cooperation

To illustrate the magnitude of conditional cooperation, we need to account for non-compliance. We focus on the results from column (3) of Table 4, because for this sample we are confident that the subjects read the letter. However, a second form of non-compliance must still be accounted for: even if they read the letter, households are not expected to fully incorporate the feedback into their beliefs. We can use the results from the Mturk Survey to correct for this additional form of non-compliance. Because each additional 1 pp in the information shock increased the perceived average tax rate by 0.393 pp (column (1) of Table 3), we should use a scale-up factor of 2.54 ($= \frac{1}{0.393}$). Scaling-up the coefficient on $\overline{\tau}$ from column (3) of Table 4 implies that increasing a household's perception of the average tax rate paid in the county by 0.1 pp would decrease the protest probability by 3.19 pp ($= 0.1 \cdot 12.566 \cdot 2.54$). The 0.1 pp reduction in the average tax rate corresponds to a 5% reduction relative to the baseline (2.01 pp), while the 3.19 pp reduction in the protest probability corresponds to 6.3% of the average protest probability reported in column (3). These percent-changes suggest that the elasticity of the protest probability with respect to the perceived average tax rate is 1.26 ($= \frac{6.3}{5}$); in other words, households are elastic to their perceptions of the taxes paid by everyone else.

We can better illustrate the magnitude of conditional cooperation with the following counterfactual analysis. In policy A, we increase a household's tax rate by 0.1 pp but hold constant the tax rates paid by everyone else. In policy B, we increase a household's tax rate by 0.1 pp, but at the same time increase the tax rates of everyone else by 0.1 pp. We further assume that households have accurate perceptions of these changes in their own and average tax rates. We can use the coefficient on τ_i from column (3) of Table 4 to approximate the effect of policy A: the protest probability would increase by 9.63 pp. The effect of policy B can be approximated by combining the coefficient on τ_i and the scaled-up $\overline{\tau}$: the protest probability would increase by 6.44 pp (= 9.63 - 3.19); in other words, the conditional cooperation is dampening a full third ($= \frac{9.63-6.43}{9.63}$) of the effect of the tax hike.

An alternative way to assess the magnitude of the coefficient on $\overline{\tau}$ is to compare it to the corresponding change in τ_i, which would generate the same effect. Using the same scaled-up coefficients from the previous paragraph, we find that the effect of a 1 pp increase in the perceived average tax rate is equivalent to the effect of a 0.33 pp reduction in one's own tax rate. This ratio is remarkably similar to what has been found in laboratory experiments of the public good game. Take, for example, Fischbacher et al. (2001), one of the earliest studies on conditional cooperation. According to Figure 1 from that paper, for each additional dollar in the average contribution of others, the average subject is willing to contribute an additional 0.35 dollars.[45] This ratio of 0.35 is close to the 0.33 ratio reported above.

5.7 Heterogeneity by Partisanship

One of the most robust findings from laboratory games is that individuals are heterogeneous in their degree of conditional cooperation. While some subjects are willing to match one-to-one the contributions made by others, other subjects do not care about the contributions of others at all (Gächter, 2007). It is then fitting to explore whether this result holds outside the laboratory as well. Given the survey-based evidence that Democrats and Republicans are rather divided in their preferences for redistribution, exploring heterogeneity by partisan identity seems like a natural starting point.

To split the analysis by Republican and Democrat households, we matched the individual-level data from the property tax records to the Texas voter files. We constructed a proxy for whether each taxpayer is likely to be a Democrat or a Republican. The details of this categorization as well as a validation exercise are in Appendix A.13.[46] 55% of the subjects are categorized as more likely to identify as Democrats and the remaining 45% are categorized as more likely to be Republican. Indeed, this narrow advantage for the Democrats is consistent with recent electoral results from Dallas County: for example, in the 2012 presidential election, Barack Obama received 57% of the votes while Mitt Romney received 42% of the

[45] Unfortunately, the authors do not report the slope from Figure 1, so 0.35 is our best guess.

[46] A majority of the sample can be categorized as Democrat or Republican directly using their participation in primary elections. We then use a simple predictive model to extrapolate to the rest of the sample. This prediction model is complemented by other data sources, such as precinct-level data on presidential voting results.

votes.

The last two columns of Table 4 break down the baseline results from column (1) by Republican and Democrat subjects. Column (7) of Table 4 presents the estimates for the Democrat subjects, while column (8) corresponds to the Republican subjects. First, the coefficients on $\overline{\tau}$ suggest that the conditional cooperation preferences are driven primarily by Democrats: the coefficient for Democrats is large (-1.325) and statistically significant (p-value=0.030), but the coefficient for the Republicans is close to zero (-0.253) and statistically insignificant (p-value=0.778). This finding suggests that there are indeed partisan differences in the degree of conditional cooperation; however, notice that the difference between the two coefficients is not precisely estimated and is thus statistically insignificant (p-value=0.323).

We can also compare the partisan heterogeneity in the role of private costs and private benefits. While the coefficients on *Extra Aid Message* are statistically significant for both Democrats and Republicans, the magnitude is more than twice as large for Republicans (2.480) as for Democrats (1.076), and their difference is highly statistically significant (p-value=0.012). Likewise, the coefficient on τ_i are statistically significant for both Democrats and Republicans, but the magnitude is 29% larger for Republicans than for Democrats. This difference is statistically significant (p-value=0.071). These two findings suggest that selfish motives are important for both Republicans and Democrats, although somewhat more important for Republicans. These partisan differences in behavior are qualitatively consistent with documented partisan differences in survey data (Stantcheva, 2020). However, in quantitative terms the partisan differences that we find are not as dramatic as some of the differences documented in the survey data.

We note that these partisan differences may be attributed to underlying differences in values and beliefs between Republican and Democrats such as trust in government or altruism. The partisan differences may also be due to other differences between Republican and Democrats, such as their demographics (e.g., income, age) or the probability of having kids enrolled in public schools.

6 Conclusions

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

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

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

More importantly, our experimental findings hints at potential solutions to these problems. Because households are so responsive to aid, access to protests may be made more equal by targeting aid to disadvantaged groups (Finkelstein and Notowidigdo, 2019). For example, by one account presented in Section 2, Hispanic households were 1.97 pp less likely to protest than comparable White households. In comparison, we find that our letter with extra aid increased the protest rate of Hispanic households by 2.53 pp (p-value<0.001).[47] Thus, sending the extra aid letter just to the Hispanic households would have been more

[47] This result corresponds to the regression from column (4) of Table 2, but estimated with the 20.39% of the sample that was classified as Hispanic.

than sufficient to close the gap between Hispanics and Whites. Some low-cost interventions of this type – promoted internally by county assessor's offices or externally through NGOs – could go a long way into making the system to protest taxes more equitable.

References

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

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

Avenancio-León, C. and T. Howard (2019). The Assessment Gap: Racial Inequalities in Property Taxation. *Available at SSRN 3465010*.

Benzarti, Y. How Taxing is Tax Filing? Using Revealed Preferences to Estimate Compliance Costs.

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.

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

Bottan, N. and R. Perez-Truglia (2017). Choosing Your Pond: Location Choices and Relative Income. *NBER Working Paper No. 23615*.

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

Cabral, M. and C. Hoxby (2012). The Hated Property Tax: Salience, Tax Rates, and Tax Revolts. Technical Report 18514, National Bureau of Economic Research.

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

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

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

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

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

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

Epper, T., E. Fehr, and J. Senn (2020). Other-regarding preferences and redistributive politics. Technical Report 339.

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

Fischbacher, U., S. Gächter, and E. Fehr (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters 71*(3), 397–404.

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

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

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

Gächter, S. (2007). Conditional cooperation: Behavioral regularities from the lab and the field and their policy implications. In *Economics and psychology: A promising new cross-disciplinary field.*, CESifo seminar series., pp. 19–50. Cambridge, MA, US: MIT Press.

Group, T. W. B. (2019). The Administrative Review Process for Tax Disputes : Tax Objections and Appeals in Latin America and the Caribbean - A Toolkit.

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

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

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

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

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.

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

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

Slemrod, J. (2018). Tax Compliance and Enforcement. *Journal of Eonomic Literature Forthcomin.*

Stantcheva, S. (2020). Understanding Tax Policy: How Do People Reason? Technical report, National Bureau of Economic Research.

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

Figure 1: Descriptive Statistics about Tax Protests

a. By Tax Amount **b.** By Ethnicity

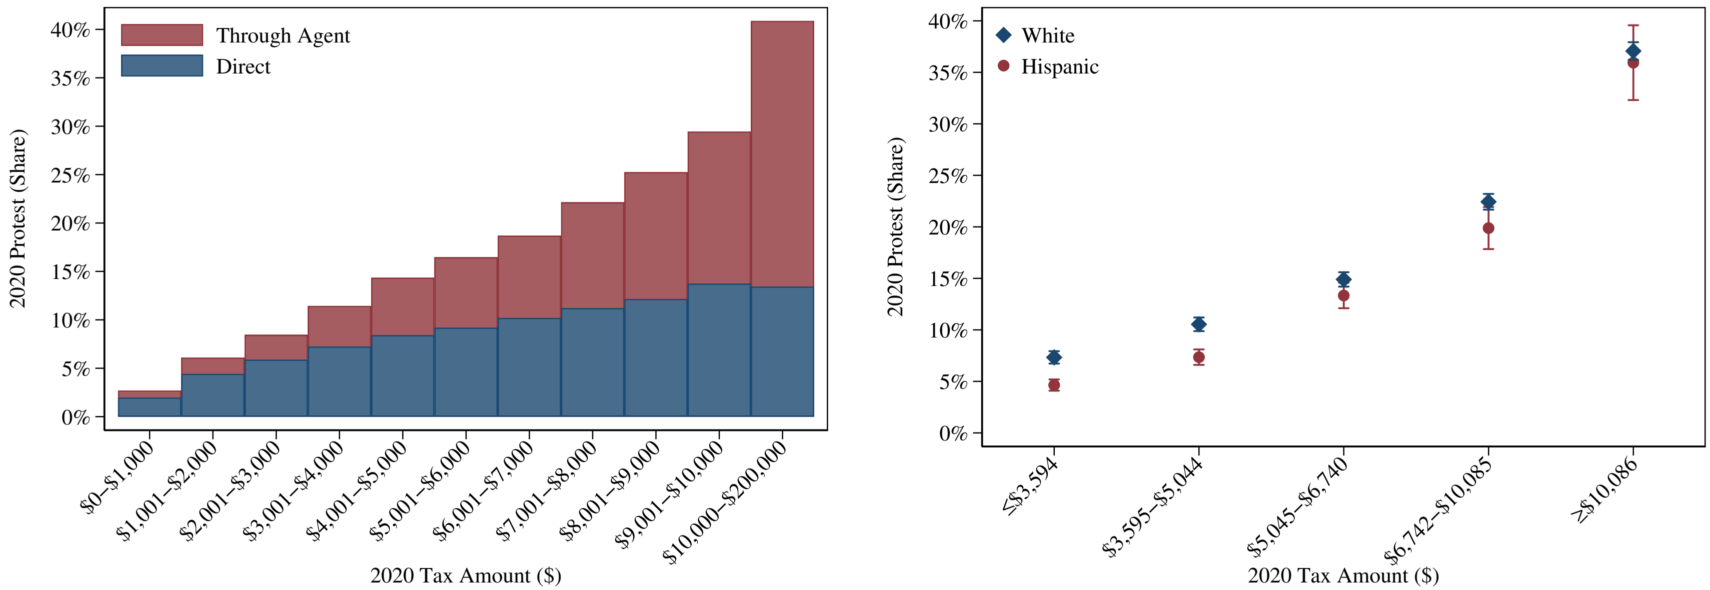

Notes: This figure presents descriptive statistics based on the sample of 423,607 single-family homes who were subject to property taxes in Dallas County in 2020. Panel (a) breaks the sample down by the estimated taxes based on the proposed values announced on May 15th, 2020. The height of the bars represent the overall share in each group that protested in 2020. The stacked bars differentiate between the households who protested through agents (in red) or directly (in blue). Panel (b) is based on a subsample of 45,511 households for which we obtained data on ethnicity and it indicated that they were either Hispanic or White. The x-axis denotes the five quintiles of the estimated taxes based on the proposed values announced on May 15th, 2020. Each dot represents the share of those households who protested directly in 2020, with 90% confidence intervals in parentheses. Red dots correspond to the households identified as Hispanic, while blue dots correspond to White households.

Figure 2: Effects of the Homestead Cap on the Probability of Protesting

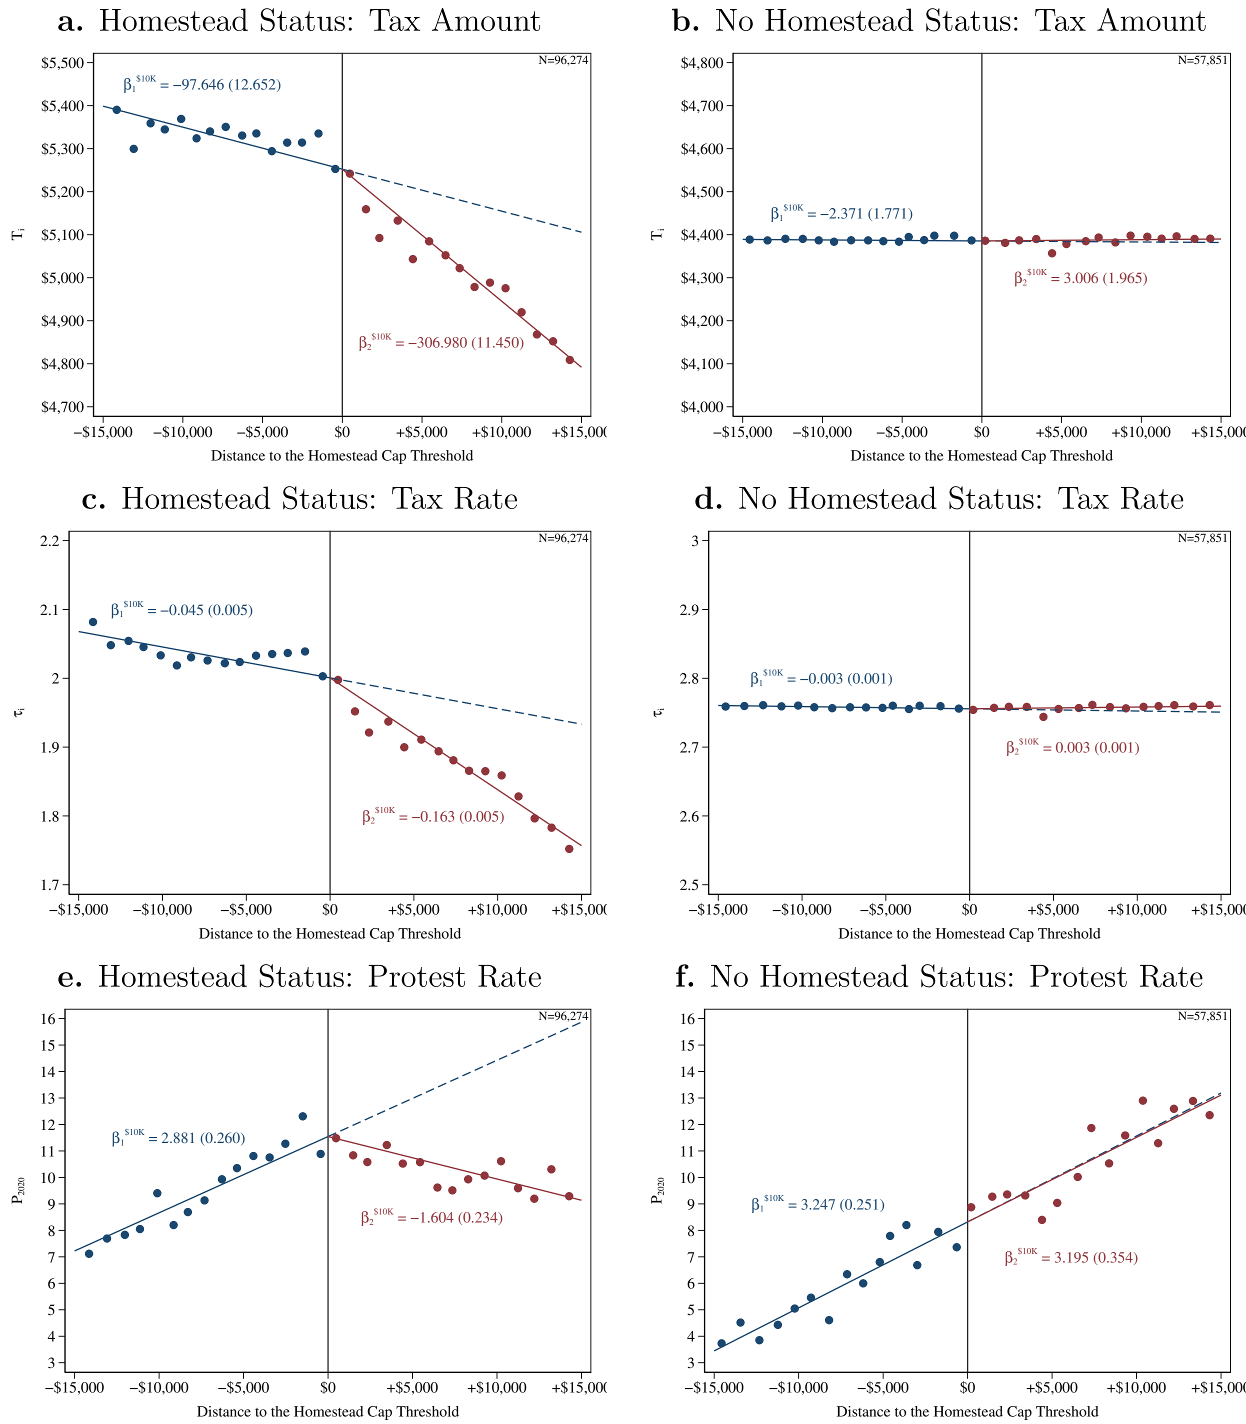

Notes: This figure features a binned scatterplot of the relationship between a given outcome (indicated in the y-axis of each panel) and the distance between the 2020 proposed value and the 2020 potential homestead cap threshold (defined as 110% of the appraised value in the previous year). All regressions 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 as robust standard errors (in parentheses) and the number of households (in brackets). The panels on the left half ((a) (c) and (e)) corresponds to households with 2020 homestead status, while the panels on the right half ((b) (d) and (f)) correspond to households without 2020 homestead status. The dependent variables are: T_i is the estimated 2020 tax amount based on 2020 proposed values; τ_i is the corresponding tax rate; P_{2020} is an indicator variable that takes the value 100 if the household protested directly in 2020 and 0 otherwise.

Figure 3: First Page of Sample Letter

THE UNIVERSITY OF TEXAS AT DALLAS
Naveen Jindal School of Management

May 15th, 2020

Dear Joan Robinson,

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

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

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

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

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

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

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

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

800 W. Campbell Road
Richardson, TX 75080

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

Please recycle

43137

Notes: A sample of the first page of the letter used in the field experiment. The information in the table varied by treatment group. Sample tables for every treatment group are presented in Figure A.6. 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 Sample Letter

a. Extra Aid Message: No

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

Thank you for your attention!

Alejandro Zentner
Associate Professor
University of Texas at Dallas

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

b. Extra Aid Message: Yes

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

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

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

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

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

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

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

Thank you for your attention!

Alejandro Zentner
Associate Professor
University of Texas at Dallas

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

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

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

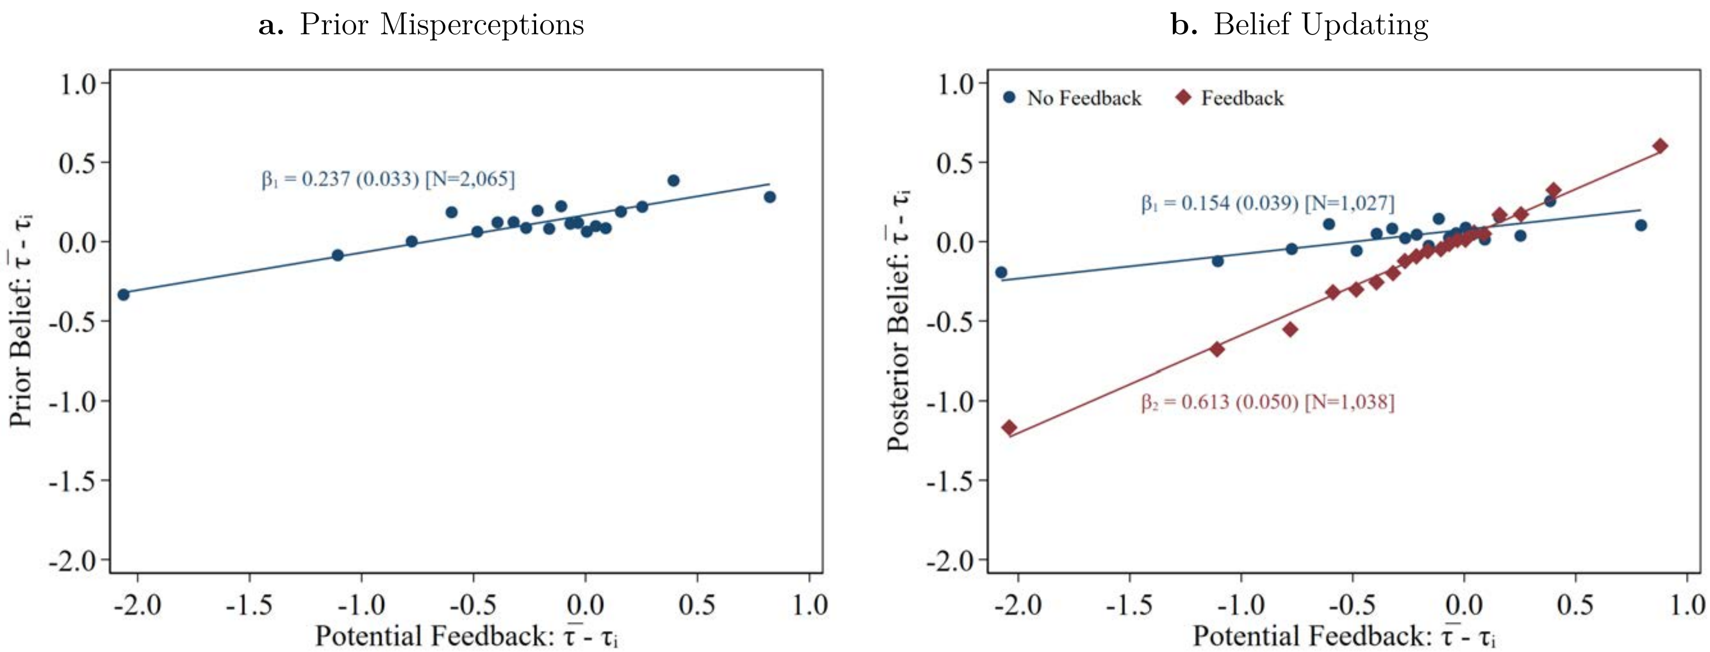

Notes: This figure shows binned scatterplots based on the Mturk Survey. Each line corresponds to a separate OLS regression, with robust standard errors in parentheses and the number of observations in brackets. $\overline{\tau} - \tau_i$ refers to the difference between the average tax rate in the respondent's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

Table 1: Effects of the Homestead Cap on the Probability of Protesting

	(1) P_{2020}	(2) P_{2019}	(3) P_{2020}	(4) P_{2019}	(5) P_{2020}^{agent}	(6) P_{2020}^{all}	(7) P_{2020}	(8) P_{2020}
τ_i	36.507***	-5.740			6.174**	42.681***	30.644***	20.928***
	(4.979)	(3.611)			(2.954)	(5.506)	(1.436)	(0.526)
T_i			0.022***	-0.004				
			(0.004)	(0.002)				
Bandwidth	\$15K	\$15K	\$15K	\$15K	\$15K	\$15K	\$30K	\$150K
Mean Outcome	9.98	6.12	9.98	6.12	5.93	15.91	8.94	8.21
Std. Dev. Outcome	29.97	23.97	29.97	23.97	23.62	36.58	28.53	27.45
Observations	96,274	96,274	96,274	96,274	96,274	96,274	179,452	308,000

Notes: Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. Each column presents results from a different Instrumental Variable regression that follows the specification presented in equation (6) from Section 3.2. The endogenous variable (τ_i) corresponds to the property tax rate (in percentage points) that the household is subject to, except in columns (3) and (4) where it is the tax amount (in dollars, T_i). The excluded instrument is the interaction between two variables: the difference between the proposed value and the potential homestead cap, and an indicator variable for whether that difference is positive. The regression controls for the difference between the proposed value and the potential homestead cap as well as a host of additional variables – see Section 3.2 for the full list. Results based on households with single-family homes and 2020 homestead status. The row *Bandwidth* indicates how the sample has been further restricted based on the absolute difference between the proposed value and the potential homestead cap. 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_{2019} indicates a direct protest in 2019; P_{2019}^{agent} indicates a protest through an agent in 2020; P_{2019}^{all} indicates any type of protest (direct or through agent).

Table 2: Effects of the Two Types of Letters on the Probability of Protesting

	(1) P_{2020}	(2) P_{2019}	(3) P_{2020}^{agent}	(4) P_{2020}^{all}	(5) P_{2020}	(6) P_{2020}	(7) P_{2020}^{online}	(8) SO_{2020}	(9) P_{2020}^{won}	(10) Δ_{2020}
Basic Aid Letter$^{(i)}$	1.792***	-0.286	0.030	1.822***	1.449***	1.935***	1.591***	0.795	0.939***	0.065***
	(0.249)	(0.199)	(0.181)	(0.283)	(0.347)	(0.339)	(0.231)	(0.719)	(0.196)	(0.018)
Extra Aid Letter$^{(ii)}$	3.509***	-0.288	-0.122	3.387***	3.108***	3.745***	3.306***	15.287***	1.939***	0.094***
	(0.258)	(0.198)	(0.179)	(0.288)	(0.364)	(0.350)	(0.241)	(0.979)	(0.203)	(0.017)
P-value (i)=(ii)	<0.001	0.995	0.409	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001	0.125
CAD Notification					No	Yes				
Mean Outcome (No Letter)	8.67	6.14	6.07	14.74	6.03	10.33	7.18	3.37	5.07	0.41
Std. Dev. Outcome (No Letter)	28.14	24.00	23.89	35.45	23.80	30.43	25.82	18.05	21.94	1.95
Observations	78,462	78,462	78,462	78,462	30,356	48,106	78,462	5,026	78,462	78,462

Notes: Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. Each column presents results from a different regression with two 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. Columns (5) and (6) split the sample by whether DCAD mailed the subjects an official notification. 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_{2019} indicates a direct protest in 2019; P_{2019}^{agent} indicates a protest through an agent in 2020; P_{2019}^{all} indicates any type of protest (direct or through agent); P_{2020}^{online} indicates a direct protest that was filed online; SO_{2020} is defined for a 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. P_{2020}^{won} indicates with 100 if a direct protest resulted in a reduction in the assessed value. And Δ_{2020}^{all} is the percentage point reduction in the assessed value, which by construction takes the value 0 if the household did not protest or if the protest was unsuccessful.

Table 3: Effects of the Information on Average Tax Rate in the County on Beliefs and Attitudes in the Mturk Survey

	(1) $\overline{\tau}_i^{post}$	(2) $\overline{\tau}_i^{prior}$	(3) τ_i	(4) R_i^{post}	(5) H_i^{post}	(6) $\Delta\tau_i^{post}$
$\overline{\tau}$	0.393***	0.028	-0.066	-0.319*	-5.135	-0.138**
	(0.071)	(0.086)	(0.044)	(0.169)	(4.052)	(0.067)
Mean Outcome	1.24	1.39	1.28	6.14	42.57	0.37
Std. Dev. Outcome	0.77	0.92	0.80	2.03	49.46	0.57
Observations	2,065	2,065	2,065	2,065	2,065	2,065

Notes: Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. Each column presents results from a different regression using data from the Mturk Survey. All regressions follow the same specification given by equation (9) from Section 5.4. The coefficient on $\overline{\tau}$ corresponds to the information shock ($D_i \cdot (\overline{\tau} - \tau_i)$). The dependent variables are defined as follows: $\overline{\tau}_i^{post}$ is the posterior belief (i.e., after the information-provision experiment) on the average tax rate in the county. $\overline{\tau}_i^{prior}$ is the corresponding prior belief (i.e., before the information-provision experiment). τ_i is the posterior belief about the respondent's own tax rate. R_i^{post}, elicited after the information-provision experiment, measures if the individual thinks that his or her own taxes are higher than the ones paid by others in the county, in a 1-10 scale. H_i^{post}, elicited after the information-provision experiment, is an indicator taking the value 100 if the respondent thinks that his or her own taxes are too high. $\Delta\tau_i^{post}$ is the desired tax cut: i.e., the difference between the tax rate that the household currently pays and the tax rate that the household would find most fair (both measured after the information-provision experiment).

Table 4: Effects of the Information on Average Tax Rate in the County on the Probability of Protesting

	All		Survey Respondents				All	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	P_{2020}	P_{2019}	P_{2020}	P_{2019}	U_{2020}	I_{2020}	P_{2020}	P_{2020}
Field Experiment:								
$\overline{\tau}$	-0.937*	-0.398	-12.566**	-2.284	-0.459*	-11.919***	-1.325**	-0.253
	(0.509)	(0.340)	(5.424)	(2.822)	(0.244)	(4.495)	(0.611)	(0.897)
Extra Aid Message	1.719***	0.013	27.373***	0.821	0.083	12.795***	1.076***	2.480***
	(0.274)	(0.202)	(2.206)	(1.355)	(0.098)	(1.758)	(0.349)	(0.433)
Quasi-Experiment:								
τ_i	38.158***	-2.103	96.314***	-21.086	1.730**	56.739***	36.992***	47.827***
	(2.270)	(1.646)	(24.284)	(14.663)	(0.867)	(19.575)	(2.927)	(5.230)
Subsample							Dem.	Rep.
Mean Outcome	11.29	5.85	50.26	10.33	7.12	81.90	9.91	12.97
Std. Dev. Outcome	31.65	23.48	50.01	30.44	2.15	38.52	29.88	33.59
Observations	50,394	50,394	1,888	1,888	1,888	1,867	27,633	22,761

Notes: Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. Each column presents results from two sets of regressions based on the same sample and with the same dependent variable. All results are based on households who received a letter in the field experiment. The estimates from *Field-Experiment* were estimated based on equation (10) from Section 5.5. The coefficient on $\overline{\tau}$ corresponds to the information shock $(D_i \cdot (\overline{\tau} - \tau_i))$, while *Extra Aid Message* corresponds to an indicator variable that takes the value 1 if the household was assigned to the extra aid message. *Quasi-Experiment* corresponds to the Instrumental Variable regression given by equation (6) from Section 3.2, in which the tax rate (τ_i, in percentage points) is the endogenous variable. This regression exploits variation in the proposed value within $150,000 of the potential homestead threshold and for the subset of households with a 2020 homestead status. The dependent variables are defined as follows: P_{2020} is an indicator variable that takes the value 100 if the owner protested directly in 2020 and 0 otherwise; P_{2019} indicates whether the owner protested directly in 2019. Columns (3) through (6) are restricted to the respondents to the Field Survey. 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. 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. Columns (7) and (8) reproduce column (1) for the subsample of Democrats and Republicans, respectively.