Journal of Accounting and Economics 81 (2026) 101838

Contents lists available at ScienceDirect

Journal of Accounting and Economics

journal homepage: www.journals.elsevier.com/journal-of-accounting-and-economics

Regular Manuscript

Paying your fair share: Perceived fairness and tax compliance☆

Brad Nathan [a,*], Ricardo Perez-Truglia [b], Alejandro Zentner [c]

[a] Rutgers University, USA
[b] University of California, Los Angeles, USA
[c] University of Texas at Dallas, USA

ARTICLE INFO

JEL classification:
C93
H2
H4
H7
Z13

Keywords:
Property taxes
Information disclosure
Tax compliance
Fairness

ABSTRACT

We provide evidence on the role of perceived fairness in tax compliance. Are households more willing to pay taxes when they believe others contribute their fair share? We investigate this question with a natural field experiment in the context of U.S. property taxes. Using an information-disclosure experiment, we exogenously shifted households' perceptions of the average tax rate paid by others. We find that higher perceived average tax rates increase perceptions of fairness and reduce the likelihood of tax appeals. For every additional $1 paid by the average household, a taxpayer is willing to contribute $0.43 more. In the field experiment, subjects were informed about the average tax rate but not why it might differ from theirs. A complementary survey shows this context matters: when households learn others pay lower rates due to exemptions such as disability or advanced age, they are more willing to tolerate unequal rates.

1. Introduction

Why do individuals pay their taxes? A natural explanation is that tax compliance stems from self-interest. For example, taxpayers may weigh the financial benefits of underreporting income against the expected penalties if caught (Allingham and Sandmo, 1972). Beyond self-interest, tax compliance may also be motivated by tax morale—that is, individuals may choose to pay taxes because they believe it is the right thing to do (Luttmer and Singhal, 2014; Slemrod, 2019). In this paper, we present novel evidence that fairness concerns can play a significant role in tax compliance. Specifically, we test the hypothesis that households are more willing to pay taxes when they believe others are contributing their fair share.[1] We examine the causal effect of fairness perceptions on

☆ Portions of this paper were originally part of a longer paper that included analyses from another treatment arm, first posted in September 2020 as NBER Working Paper No. 27816. As the original paper became too lengthy, it was split into this paper and Nathan et al. (2025). We are grateful for valuable feedback from Sudiptha Basu, Raj Chetty, Christine Cuny, Austan Goolsbee, Steve Levitt, Ningzhong Li, John List, Joel Slemrod, Erzo Luttmer, Anya Nakhmurina, James Poterba, Gurpal Sran, and seminar participants at the NBER-Public Economics, Stanford University, University of Chicago, Yale University, Rutgers Accounting Research Conference, Federal Reserve Bank of Chicago, AEA, NTA, ZEW, and AEI. This project was reviewed and approved in advance by the Institutional Review Board at The University of Texas at Dallas. The experiments were registered in the AEA RCT Registry (#0005992). The original pre-registration was posted on May 24, 2020. However, after receiving interview requests from the media, we removed the pre-registration (we were concerned that the media would divulge the hypotheses listed in the pre-registration and thus contaminate the study). We re-posted the pre-registration on June 16, 2020, after the deadline to file an appeal had passed (and before the data to conduct the analysis became available). Adrian Cadena Medina, Luisa Cefala, Dongwook Chun, Karl Dill, Santiago De Martini, and Xinmei Yang provided superb research assistance.

* Corresponding author.
E-mail addresses: brad.nathan@business.rutgers.edu (B. Nathan), ricardo.truglia@anderson.ucla.edu (R. Perez-Truglia), azentner@utdallas.edu (A. Zentner).

[1] Our hypothesis relates to concepts such as reciprocal fairness (Luttmer and Singhal, 2014) and conditional cooperation (Fischbacher et al., 2001; Gachter, 2007).

https://doi.org/10.1016/j.jacceco.2025.101838
Received 3 June 2024; Received in revised form 26 August 2025; Accepted 28 September 2025
Available online 19 November 2025

households' willingness to pay taxes through a pre-registered natural field experiment using revealed preference data—specifically, households' decisions to file tax appeals.

The setting for our study is property taxes, a significant source of government revenue in the United States and around the world. In 2019, property taxes were the second-largest source of tax revenue in the U.S., generating an estimated $577 billion (Mayo et al., 2021). This figure was more than twice the revenue generated by the corporate income tax, which totaled about $230 billion (Congressional Budget Office, 2020).[2] Our field experiment took place in Dallas County, the second-largest county in Texas, with a population of 2.6 million—larger than 15 of the 50 U.S. states (US Census Bureau, 2021). Property tax systems operate nearly identically across all 254 counties in Texas and function similarly in other states and countries (Dobay et al., 2019; World Bank, 2019; Nathan et al., 2025). The average household in Dallas County paid $5,916 in property taxes and the average home value in the county was approximately $294,846, making it a high-stakes context for studying tax compliance. For comparison, in 2020 the average federal income tax bill was $13,704 and the average household income in Dallas County was $82,198 (IRS, 2020). Texas does not have a state income tax, so property taxes serve as a key source of revenue for funding local government services.

To test our fairness hypothesis, the ideal experiment would hold a household's own tax rate constant while randomly varying the average tax rate paid by other households. According to our hypothesis, conditional on a household's own tax rate, we would expect greater willingness to pay taxes when the average tax rate is higher. In other words, while households may always have an incentive to free-ride on the taxes paid by their neighbors, they may feel a stronger moral obligation to pay their own taxes if they believe others are contributing their fair share. Moreover, when individuals learn that others are paying less, their reaction may depend on what they perceive to be the reason for that difference. If someone finds that their neighbors pay a lower tax rate but receives no additional context, this may increase perceived unfairness and, in turn, reduce tax compliance. However, if the individual believes other households face lower tax rates due to extenuating circumstances—such as disability, old age, or lower income—they may view the disparity as justified, thus mitigating or eliminating the effects on perceived unfairness and tax morale. Therefore, in the ideal experiment, we would also randomize the reason why some households face lower tax rates, such as due to arbitrary reasons or extenuating circumstances. In practice, however, such experiments are not feasible, as randomizing tax rates would be impractical, potentially illegal, and unethical. Instead, our research design introduces randomized shocks to households' perceptions of the average tax rate in the county through an information-disclosure experiment.

In this paper, we present results from a field experiment in which some subjects were randomly provided with information about the average tax rate in their county. We find that when households learn that others pay a lower tax rate, their perceptions of unfairness increase, and they become more likely to file a tax appeal. In this field experiment, no additional context was provided about why the average tax rate might differ from their own. Additionally, we use a complementary survey (the *Preferences Survey*) to show that context matters: on average, households perceive less unfairness when others pay a lower tax rate due to extenuating circumstances such as a disability or advanced age.

Our setting offers two key advantages for studying tax fairness. First, there is substantial variation in effective property tax rates—hereafter referred to simply as tax rates, defined as the total property tax amount divided by the home's assessed value. The average household in the county faces an effective tax rate of 2.01%, but some households pay significantly more or less. For example, the 10th percentile of the tax rate distribution is 1.33%, while the 90th percentile is 2.73%. Because the tax schedule is proportional, these differences are not driven by variation in assessed property values. Instead, we show that most of the variation arises from exemptions and caps granted to specific groups, such as primary residence owners, individuals with disabilities, and older adults.

The second advantage of our setting is that, despite the high-enforcement nature of property taxes, we can still measure willingness to pay taxes through revealed preferences—specifically, by observing the decision to file a tax appeal.[3] Filing an appeal is a consequential action that households can take to legally reduce their tax burden. The appeal process operates as follows: each year, the Dallas Central Appraisal District (DCAD) determines a proposed assessment of a property's market value, referred to as the proposed value. Property taxes are then calculated based on this assessment. Because the appraisal process involves considerable ambiguity and subjectivity, households have the opportunity to contest the county's assessed value of their homes by filing a property tax appeal—also known as a tax protest. Many households take advantage of this process, often securing substantial reductions in their tax bills. For example, among households in our field experiment, 11.3% filed an appeal. Of these appeals, 76.6% resulted in a reduction in the assessed value. Successful appeals yielded an average savings of $457 (7.7% of the average property tax bill) in the first year alone.

In the first part of the paper, we present results from the *Perceptions Survey*, which validates the design of the field experiment. We recruited subjects via an online platform and began by eliciting their beliefs of the average tax rate in their county. Next, a randomly selected half of the sample received information about the actual average tax rate. Finally, we re-elicited participants' beliefs. The results reveal that homeowners hold substantial misperceptions about the county's average tax rate. In particular, households tend to assume that the average tax rate is similar to their own. As a result, those with below-average tax rates tend to underestimate the true average, while those with above-average rates tend to overestimate it. When presented with accurate information, subjects update their beliefs in a predictable manner: households with below-average tax rates revise their beliefs upward, while those with above-average rates revise them downward. This systematic pattern of belief updating forms the foundation of the identification strategy used in the field experiment.

[2] For reference, the federal individual income tax generated $1.7 trillion in 2019 (Sherlock and Marples, 2021).

[3] For other studies leveraging property tax appeals, see Jones (2019), Nathan et al. (2025), and Giaccobasso et al. (2025). For examples of field experiments in accounting and economics, see Floyd and List (2016).

The field experiment involved sending letters to a sample of 50,394 households in Dallas County at the start of the period during which they could file a tax appeal. All letters included information about the household's own property taxes. In addition, we randomized whether the letter also included accurate information about the taxes paid by the average household in the county. This randomization was designed to shift beliefs in line with the patterns observed in the Perceptions Survey: households with below-average tax rates were expected to revise their beliefs upward about the average, while those with above-average rates were expected to revise them downward. We measure whether these information shocks had a causal effect on homeowners' perceptions of fairness, using survey data, and on their willingness to file a tax appeal, using administrative records.

In an additional survey—hereafter referred to as the *Field Survey*—we collected responses from households included in the field experiment. The letters sent to participants contained a URL linking to a brief survey, along with a unique survey code that allowed us to identify individual respondents. We leverage these survey responses in two key ways. First, we use a household's completion of the Field Survey as an indicator of whether it paid close attention to the letter. Second, we included a specific question on perceived fairness to serve as a dependent variable in the experimental analysis. Specifically, we asked respondents whether they perceived their own taxes as fair or unfair compared to the average household.

The results from the field experiment indicate that the information shocks significantly influenced real-world behavior in a manner consistent with the fairness hypothesis. When a household learns that the average tax rate is higher (lower) than it initially believed, this decreases (increases) the probability of filing a tax appeal. It is important to note that this is an intention-to-treat estimate. As widely documented in mailing experiments (e.g., Perez-Truglia and Cruces, 2017; Gerber et al., 2020; Bottan and Perez-Truglia, 2025), a large fraction of households may not have opened the letter or may not have paid close attention to its content. To mitigate this attenuation bias, we focus on the subsample of 1,888 respondents to the Field Survey, for whom we have confirmation that they engaged with the letter's content. As expected, the information shocks have substantially larger effects within this group. Moreover, using the question on perceived fairness from the Field Survey, we can test whether the effects of the information shocks on tax appeals operate through the fairness channel. We find that, consistent with the fairness channel, when households learn the average tax rate is higher (lower) than they initially believed, they are significantly more likely to perceive their own taxes as more fair (less fair).

To assess the magnitude of this effect on tax appeals, we correct for a second source of non-compliance common to information-provision experiments: when subjects receive new information, they tend to revise their beliefs in the expected direction, but not fully (see, e.g., Bottan and Perez-Truglia, 2025). We use the degree of belief updating observed in the Perceptions Survey to adjust for this partial compliance. The resulting estimate suggests that a 0.1 pp increase in the perceived average tax rate reduces the likelihood of filing a tax appeal by 2.75 percentage points (pp). Moreover, we can quantify this fairness channel using a money metric. To do so, we draw on quasi-experimental evidence from Nathan et al. (2025), which estimates the effect of expected tax savings on the probability of filing a tax appeal. These calculations suggest that for every additional $1 paid by the average household, a household is, on average, willing to contribute an additional $0.43 in taxes. Remarkably, this magnitude closely aligns with findings from laboratory public goods experiments, where an additional $1 in average contributions increases an individual's own contribution by $0.45 (Fischbacher et al., 2001).

In addition, we present results from the Preferences Survey, based on a sample of 528 American homeowners. In the field experiment, we provided information about the tax rate paid by the average household but did not explain why this rate might be higher or lower than that of the recipient. Certain features of the tax system can lead to differences in households' tax rates that may be rooted in principles of fairness. For example, exemptions provide tax breaks for primary residences, older adults, and individuals with disabilities. While taxpayers may view it as unfair that another household pays a lower rate, they might react less negatively if they believe, for instance, that the household benefits from a disability exemption. The Preferences Survey was designed to test these additional hypotheses.

The survey begins by presenting respondents with a pair of households and asking them to set property tax rates for each, subject to a fixed tax revenue target. Depending on the treatment group, the two households could be identical or differ in specific ways. For example, in one treatment, one household included an individual with a disability. The results show that most respondents set unequal tax rates when one household had a special circumstance such as a disability or older age. Still, 30% to 45% preferred to equalize tax rates across households, even when one household had a special circumstance.

In a second survey module, respondents were shown two households facing different property tax rates and were asked whether they perceived this inequality as fair or unfair. In the control group, the households were identical, while in the treatment groups, the household paying the lower rate had a distinguishing characteristic, such as an owner with a disability. The results from this module reinforce the findings from the first. For example, when identical households pay different rates, the average perceived fairness score is 2.3 out of 10. In contrast, when the lower-tax household has a special circumstance (e.g., disability or old age), the perceived fairness score is still far from perfect but more than twice as high: 5.8 out of 10 on average.

The results from the Preferences Survey have important implications for interpreting the findings from the field experiment. In the field experiment, we provided information about the average tax rate but did not offer any context for why some households pay more than others. Some participants may have correctly inferred that the variation was due to factors such as primary residence or age-based exemptions. However, others may not have been able to make this inference on their own. If all participants correctly inferred the reasons behind the variation, then providing this additional context would not have changed our interpretation of the findings from the field experiment. But if some portion of participants failed to make this inference, the interpretation of our results would shift. In the extreme case—if no participants inferred the context on their own—the Preferences Survey suggests that, had they been aware of the full context, their reactions to the average tax rate information would have probably been about half as strong.

Our study contributes to the literature on tax morale (Luttmer and Singhal, 2014). The accounting and economics literature has documented that taxpayer behavior aligns with the predictions of cost-benefit analyses, such as research showing that higher audit rates or harsher penalties reduce evasion (Hoopes et al., 2012, 2018; Bergolo et al., 2018; Baugh et al., 2018; Belnap et al., 2023).[4] However, there is limited compelling evidence to support the role of tax morale (Slemrod, 2019; Giaccobasso et al., 2025). There are a few exceptions. Most related to our experiment, Hallsworth et al. (2017), Del Carpio (2022), and Tanner and Manwaring (2023) show evidence suggesting that providing information on descriptive norms (e.g., a message that "9 out of 10 households pay their taxes on time") can reduce late payments and tax delinquency. There is also a body of literature studying the role of fairness in tax compliance using laboratory experiments and survey data. For example, using a laboratory experiment, Alm et al. (2017) show that individuals are more likely to comply with their tax obligations, such as filing and reporting their taxes accurately and refraining from cheating, when they believe that other individuals are fulfilling their obligations as well. Using survey data, Ballard-Rosa et al. (2017) show how economic and fairness concerns shape income tax preferences.[5] We contribute to this literature by providing evidence that fairness considerations are significant in the natural, real-world context of property taxes and based on revealed preferences with high-stakes behavior. This approach addresses one challenge with studying tax compliance using survey or laboratory data, which is that they may not capture how people respond when they face real stakes.[6] Finally, our findings are also consistent with, and complementary to, evidence on other normative considerations such as benefit-based taxation or reciprocity (Giaccobasso et al., 2025; Chow et al., 2023; Belnap et al., 2025).[7]

The remainder of the paper proceeds as follows. Section 2 describes the institutional context. Section 3 presents the Perceptions Survey. Section 4 discusses the design of the field experiment, while Section 5 presents its results. Section 6 presents the Preferences Survey. The last section concludes.

2. Institutional context

2.1. Dallas County

The field experiment takes place in Dallas County. The county is diverse along various dimensions, including ethnicity and partisan identity. For example, in the 2024 presidential election, Kamala Harris received 60% of the votes, while Donald Trump received 38% (Dallas County Elections Department, 2024). We use publicly available administrative data provided by the Dallas Central Appraisal District (DCAD, 2020). These data include information on the names and mailing address of the property owners, property address (which may not coincide with the mailing address of the owners), and details about the properties, such as their characteristics (e.g., number of bedrooms, square footage), taxable values, tax rates, and history of property tax appeals. Our study focuses specifically on single-family residential homes.

2.2. Property taxes

Property taxes help finance a wide range of local services, both at the county level (such as community colleges) and within smaller jurisdictions like cities and school districts. The countywide average effective property tax rate is 2.01%, although rates can vary substantially across households. The total tax bill for a household is determined by the property's assessed value, the tax rates set by each jurisdiction, and any exemptions for which the household qualifies. For instance, homeowners can designate their primary residence as a homestead, which—among other protections—caps the annual increase in assessed value at 10%.

2.3. Source of the differences in tax rates across households

A key feature of our setting is that households with similar homes can face very different effective tax rates. Understanding the sources of this heterogeneity is crucial for designing our field experiment, which leverages variation in perceived average tax rates. In this section, we quantify tax rate heterogeneity within the county and identify its main drivers.

One simple measure of heterogeneity in tax rates across households is the Mean Absolute Difference (MAD), defined as $\frac{1}{N}\sum_{i=1,\ldots,N}|\tau_i - \bar{\tau}|$. In Dallas County, the MAD is approximately 0.4, meaning that, on average, a household's tax rate deviates by 0.4 pp from the average tax rate. This heterogeneity in tax rates may result from differences in statutory tax rates across jurisdictions within the county. For example, two households might reside in different cities within the same county, each with its own distinct city tax rate. However, as shown below, differences in statutory rates are minor; most of the variation in tax rates stems from households being affected differently by exemptions and caps.

[4] More broadly, this study is related to an accounting literature on tax planning decisions (Hanlon et al., 2022).

[5] Others examine perceptions of tax fairness in relation to social and political attitudes (Tyran and Sausgruber, 2006; Kerschbamer and Muller, 2020; Stantcheva, 2021).

[6] See Marshall et al. (2024) for a meta-analysis of the tax fairness-tax compliance literature. There is also literature studying fairness, tax compliance and related questions using analytical models, such as Myles and Naylor (1996), Traxler (2010), and Alm and Torgler (2011).

[7] Giaccobasso et al. (2025) show that households are more willing to pay property taxes if they perceive that the government uses their tax dollars to provide government services that benefit them. Furthermore, Chow et al. (2023) demonstrate companies' tax planning is negatively related to the amount of ground-level ozone pollution, suggesting a tax-public goods reciprocity between firms and the government. Belnap et al. (2025) find a negative relationship between the sentiment of tax-related posts on the social media platform X and individuals' tax compliance aggregated at the local level, which is suggestive of reciprocity.

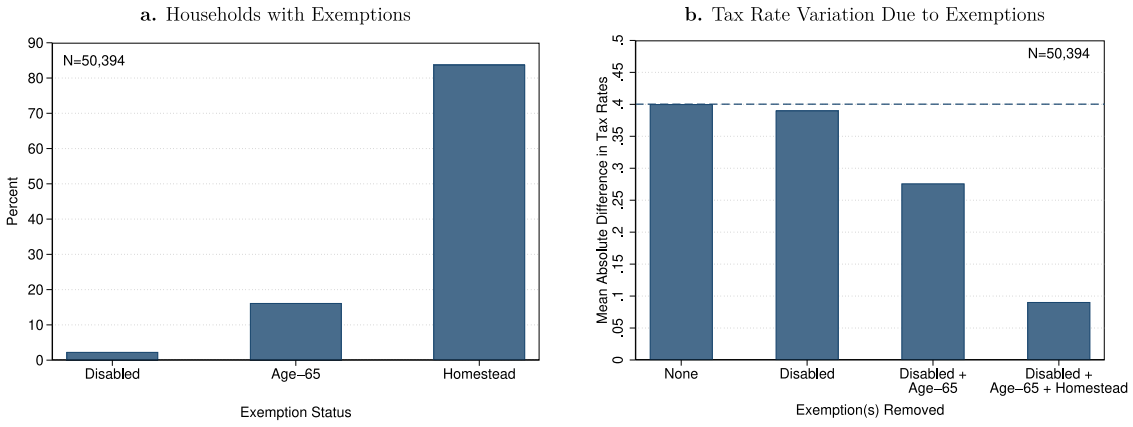

Fig. 1. Source of Variation in Tax Rates

Notes: This figure features the share of households with each exemption status and the effects of removing exemptions on the variation in tax rates for the 50,394 subjects in the field experiment who received a letter. Panel (a) shows the percentage of households that have a disability, over age 65, or homestead exemption. The vertical axis corresponds with the percentage of households that have each exemption, and the horizontal axis corresponds with each exemption status. Panel (b) shows the Mean Absolute Difference (MAD) in tax rates, defined as $\frac{1}{N}\sum_{i=1,\ldots,N}|\tau_i - \overline{\tau}|$, when removing each exemption status (and any associated caps) one at a time. The vertical axis in panel (b) represents the MAD in tax rates. The horizontal axis corresponds with four scenarios, where exemptions are removed incrementally in each scenario. In the first scenario, all exemptions are in place (i.e., no exemptions removed). In the second scenario, we remove disabled exemptions. In the third scenario, we remove the over age 65 exemptions in addition to disabled exemptions. In the fourth scenario, we remove homestead exemptions in addition to disabled and over age 65 exemptions (therefore all exemptions are removed). Tax rate is defined as the household's tax amount divided by the DCAD's proposed value for the home, then multiplied by 100. Each exemption status is defined as follows: *disabled* status is available for the primary residence of a disabled person (per the Social Security Administration) or a partly-disabled veteran (per the U.S. Veterans Administration) homeowner, or their surviving spouse; *over age 65* status is available for the primary residence of a homeowner who is age 65 or older, or their surviving spouse; *homestead* status is available for a homeowner's primary residence. See Section 2.3 for details about exemptions.

In Dallas County, homeowners may apply for three types of exemptions. The first is the disabled exemption, which applies when a homeowner has a disability and lives in the home as their primary residence.[8] The second is the over age 65 exemption, available to homeowners aged 65 or older who live in their home as their primary residence. The third is the homestead exemption, which applies when the property serves as the homeowner's primary residence.[9] These exemptions provide deductions to the home's appraised value (with amounts varying by jurisdiction), as well as tax caps.[10] The caps function differently depending on the exemption type. For households with a disabled or over age 65 exemption, participating jurisdictions (typically Independent School Districts) "freeze" the tax amount at the level paid in the year the household qualified for the exemption. The homestead cap limits annual increases in the county's appraised home value to 10% (see Nathan et al. (2025) for details).

Fig. 1(a) shows the share of households claiming each exemption status. Approximately 2.3% of households claim a disability exemption, 16.9% claim the over age 65 exemption, and 83.8% claim the homestead exemption. Even among households in the same jurisdiction with the same exemption, differences in tax rates can still arise due to how the caps function. For example, if two households claimed the over age 65 exemption, but one qualified for it recently while the other claimed it a decade ago, the latter is likely to benefit from a much larger tax savings.

Next, we examine how exemptions contribute to tax rate heterogeneity. The *y*-axis in Fig. 1(b) shows the MAD in tax rates under different counterfactual scenarios. The leftmost bar (labeled "None") represents the current tax rate distribution, with a MAD of 0.4 pp. The subsequent bars illustrate the effects of progressively eliminating exemptions and caps while holding the set of households constant. The second bar shows the impact of removing only the disability exemption, causing a slight decrease in MAD from 0.40 to 0.39. This small decrease is expected given the small share of households claiming this exemption. The third bar reflects the additional removal of the over age 65 exemption, leading to a more substantial drop in MAD to 0.28, as this exemption is more widely claimed. The fourth bar represents the elimination of the homestead exemption as well, dramatically reducing the MAD to 0.09. This sharp decline occurs because the homestead exemption is widely used, and its impact varies significantly even between households who claimed it, depending on the recent changes in home values in their area. By eliminating all three exemptions, we

[8] This includes households claiming a Disabled Person Homestead exemption (for homeowners classified as disabled by the Social Security Administration) or a partial Disabled Veteran Homestead exemption (for veteran homeowners classified as partly disabled by the U.S. Veterans Administration). Homesteads owned by fully disabled veterans are exempt from property taxes and are not included in our samples.

[9] The formal name for the homestead exemption is the "Residence Homestead" exemption, while the over age 65 exemption is known as the "Age 65 or Older Homestead" exemption. We include exemptions for surviving spouses of persons who received the Age 65 or Older Homestead or the disabled exemptions in our definitions of these respective exemptions.

[10] Appendix A.2.1 includes a brief explanation of how property taxes are calculated. For more details, see Nathan et al. (2025).

remove most of the variation in tax rates across households.[11] The little remaining variation (MAD of 0.09), by construction, reflects differences in statutory rates across jurisdictions.

These findings highlight several facets of property taxes, which are both similar and different from federal individual income taxes. Income taxes rely on progressive schedules with different tax rate brackets, while property taxes are proportional to the assessed value of the home. Both systems, however, reduce tax burdens based on taxpayer-specific circumstances. For example, property tax exemptions are available for homeowners with homestead, disability, and senior status. Similarly, income taxes allow deductions or exemptions for factors such as charitable giving, dependents, casualty losses, marital status, and more. In this sense, both tax systems incorporate elements of fairness and attempt to address perceived inequalities, albeit to limited extents (Hines, 2020).

2.4. Tax protests

Each year, homeowners have the opportunity to appeal their taxes if they disagree with their home's market value as proposed by the DCAD. Homeowners are notified of their proposed values through the DCAD's website, and some may also receive a notification by mail.[12] Homeowners can challenge the tax assessment if there are errors in the public records (e.g., incorrect square footage), if they believe that the market value of their property is lower than the proposed value, or if they believe comparable homes received lower proposed values. It is worth noting that home market values are subjective, except perhaps for homes that have recently been sold and thus have a recent transaction price. For example, a quick search on websites such as Zillow.com and Redfin.com shows wide discrepancies in market values for the same home.

Homeowners have the option of filing a tax appeal on their own, either online or by mail, which we refer to as *direct protests* following previous research (Nathan et al., 2025; Giaccobasso et al., 2025). Homeowners can also choose to hire an agent to file a protest on their behalf. These agents typically charge a combination of a flat fee and a percentage of tax savings, sometimes reaching up to 50% of the savings.[13] As stated in the pre-registration, our primary outcome of interest is direct protests. This focus is consistent with prior work (Nathan et al., 2025; Giaccobasso et al., 2025), which—for reasons explained therein—also emphasizes direct protests. However, we report results for tax agents as a secondary outcome. While households can affect their tax rate by applying for exemptions they qualify for, they cannot affect their tax rate by protesting. Specifically, when a household protests its property taxes, the protest can reduce the tax liability by lowering the assessed market value (the base), but not by reducing the tax rate (see Appendix A.2.7 for more details).

Homeowners have a one-month window from the notification date to protest their taxes. In 2020, the DCAD notified the proposed values on May 15th, and the deadline to file a protest was June 15th. Once protests are initiated, they go through a resolution process that may involve a proposed settlement by the DCAD, informal negotiations between the owner and the DCAD typically conducted online, via email, or by phone, or a formal hearing with a quasi-judicial entity known as the Appraisal Review Board. As a result of this process, the taxable value of the property is finalized (thereon called the "certified value"), unless it is further challenged in court.

3. Perceptions Survey

3.1. Survey design and implementation

This survey was designed to validate the information-disclosure experiment in the field experiment. The full survey instrument is attached as Appendix F. The survey begins by eliciting prior beliefs about the market value and property taxes of the average home in subjects' county of residence in the recent past (two years earlier, in 2018). Then, we inform all subjects that they will be randomly selected to receive or not receive accurate feedback about these values. Next, we share the feedback with those randomly selected to receive it (the treatment group), and we do not share this information with the rest of the subjects (the control group). Finally, we elicit posterior beliefs about both variables from the treatment and control groups. To avoid asking the same question twice, and following the design of other information-disclosure experiments (e.g., Cavallo et al., 2017), we elicit prior beliefs about 2018 and posterior beliefs about 2020.

This survey experiment was included as part of the same AEA RCT registration as the field experiment, and it was conducted at around the same time as the field experiment—from June 5th to June 15th, 2020. We recruited participants on Amazon Mechanical Turk, following best practices for recruiting individuals from online platforms. Participation was restricted to respondents located in the United States. In the survey announcement, we explicitly stated that the survey was limited to homeowners, and we included screening questions at the beginning of the survey. We collected responses from 2,065 U.S. homeowners. The median respondent took about 7 min to complete the survey. Additional details about the survey experiment implementation and sample are provided in Appendix A.1.

[11] Appendix A.2.2 presents histograms to further illustrate how exemptions and caps affect tax rate heterogeneity. This appendix also examines how the average tax rate changes as exemptions and caps are removed.

[12] The DCAD mails notifications to households with increased appraised values or those meeting certain criteria (for more details, refer to Nathan et al. (2025)). See Appendix E for a sample notification.

[13] For more discussion on the role of agents, see Nathan et al. (2025).

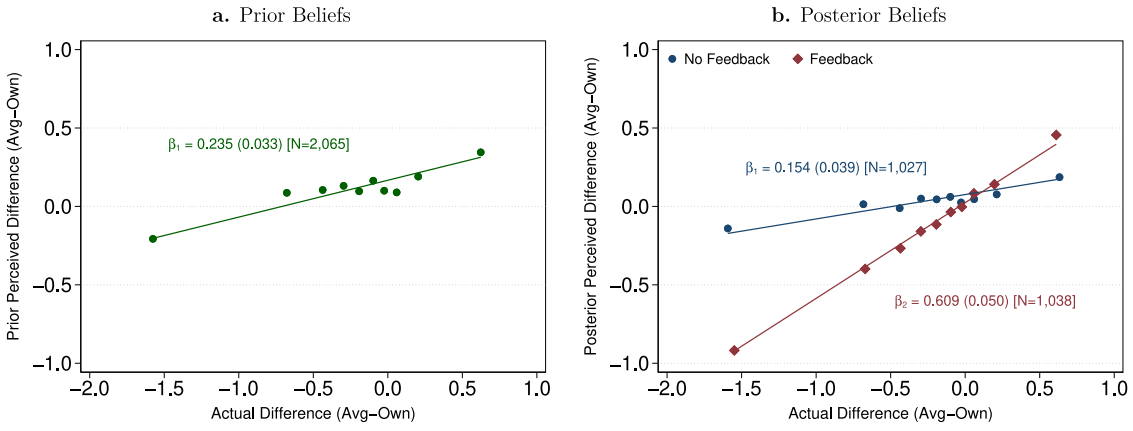

Fig. 2. Prior Misperceptions and Belief Updating in the Perceptions Survey

Notes: This figure shows how the information treatment in the Perceptions Survey affected respondents' beliefs. Each line corresponds to a separate OLS regression, with robust standard errors in parentheses and the number of observations in brackets. The x-axis corresponds to the respondent's *actual* relative position with respect to the county average (i.e., the difference between the actual average tax rate in the respondent's county and the respondent's own tax rate ($\overline{\tau}-\tau_i$)). In panel (a) the y-axis corresponds to the difference between the respondent's *perceived* average tax rate in the county and the household's own tax rate prior to the feedback on the actual county average tax rate could be shown. In panel (b) the y-axis is the respondent's corresponding *posterior* belief at the end of the survey (i.e., after the information-provision experiment). The regressions in this figure include the actual average home value and the actual average tax amount in the respondent's county as controls. The results from panel (b) are broken down by treatment group and response decile within the treatment group: the red diamonds (labeled "Feedback") correspond to deciles of respondents who were shown the feedback while the blue circles (labeled "No Feedback") correspond to deciles of those not shown the feedback.

3.2. Misperceptions about the average tax rate

The results from the survey experiment are presented in Fig. 2. In panel (a), the x-axis represents the respondent's *actual* relative tax rate: i.e., the difference between the *actual* average tax rate in the county and the household's own tax rate. The y-axis represents the respondent's *perceived* relative tax at the start of the survey: i.e., the difference between the prior belief about the average tax rate and the household's own tax rate. In a world of perfectly accurate perceptions, we would expect the observations to align precisely along the 45-degree line, corresponding to a slope of 1. At the other extreme, if subjects have no idea about the county's *actual* average tax rate, we would expect a slope of 0. We find that households' perceptions are far from accurate: the coefficient estimate (0.235) falls significantly short of 1. More precisely, Fig. 2(a) shows that misperceptions are systematically skewed toward the middle: individuals who pay more than average tend to underestimate how much they pay relative to others, and individuals who pay less than average tend to overestimate how much they pay. This type of systematic bias is expected, given that individuals have little or no information about the average rate and thus use their own tax rate as their best guess for what others pay. In fact, similar biases have been documented in other contexts. For example, when employees are asked to guess the average salary of their peers, they tend to use their own salary as the benchmark (Cullen et al., 2021). Likewise, when estimating their position in the income distribution, individuals systematically believe they are near the middle (Cruces et al., 2013).

Given the systematic biases discussed above, when accurate information is provided, we expect individuals toward the left end of the x-axis in Fig. 2(a) to update their beliefs downward, and individuals toward the right end of the x-axis to update their beliefs upward. Fig. 2(b) shows how subjects actually updated their beliefs. Just as in panel (a), the x-axis in panel (b) is the respondent's *actual* relative position with respect to the average tax rate. However, the y-axis in panel (b) is different from that in panel (a): instead of showing the prior beliefs, panel (b) presents the posterior beliefs. In other words, the y-axis represents the respondent's *perceived* relative tax rate at the end of the survey (i.e., after the information-provision experiment). The blue dots correspond to subjects in the control group (i.e., subjects who were not shown feedback about the actual average tax rate). For this group, the relationship between perceptions and reality continues to be weak, just as in panel (a). This is to be expected: since these respondents did not receive any new information, we would not expect their posterior beliefs to become any more accurate. In contrast, the red diamonds correspond to the treatment group (i.e., subjects who were shown the feedback). These red diamonds show that the correlation between perceptions and truth becomes markedly stronger when individuals are provided with accurate feedback. For a more formal test, we compare the slope between perceptions and truth in the control group (0.154) versus the corresponding slope in the treatment group (0.609). Consistent with significant learning, the difference between the two is not only large (0.455) but also highly statistically significant (p-value < 0.001).

Fig. 2(b) illustrates a key fact that motivates the experimental design for the field experiment. When individuals receive accurate information on the average tax rate, their reaction depends on their prior beliefs. For individuals with accurate prior beliefs, they should not react to the information because it does not contain any news. For individuals whose prior beliefs underestimate the average tax rate, they should react to the information by increasing their posterior beliefs. For individuals whose prior beliefs

overestimate the average tax rate, they should react to the information by lowering their posterior beliefs. In the field experiment, we do not observe the prior beliefs of the subjects. However, we can still capture the heterogeneous effects of the information, by leveraging whether the individuals have tax rates that are above or below the county average (i.e., whether the observation falls toward the left or right on the x-axis). On the one hand, individuals with tax rates above the average tend to systematically underestimate the average tax rate and, when shown the information, they update their perceptions upward. On the other hand, individuals with tax rates below the average tend to overestimate the average tax rate, so when shown the information, they will systematically update their perceptions downward. In sum, in the field experiment we can capture the heterogeneous effects of the information shock even though we do not observe prior beliefs.

4. Field experiment

4.1. Main hypothesis

Our research question can be summarized in a simple equation. Let WTP_i be household i's willingness to pay taxes—in the field experiment, we use the decision to file a tax appeal as a proxy for the (un)willingness to pay taxes. Let $\hat{\tilde{\tau}}_i$ be household i's perception about the average tax rate paid by other households in the county. We are interested in the following relationship:

$$WTP_i = \mu_0 + \mu_1 \cdot \hat{\tilde{\tau}}_i \tag{1}$$

The key hypothesis is that $\mu_1 > 0$: that is, when a household believes other households pay a higher tax rate, its own willingness to pay taxes increases—leading to a lower probability of filing an appeal. In the context of property taxes, households will always have the temptation to free-ride on the taxes paid by their neighbors. In other words, completely selfish households should be willing to pay $0 in taxes and not be affected by how much other households pay (i.e., $\mu_1 = 0$). If fairness considerations matter, households may still want to pay taxes; for example, they may consider it the right thing to do. More precisely, due to reciprocal fairness (Fehr and Schurtenberger, 2018), our hypothesis is that households feel a heightened moral obligation to pay their own taxes when they perceive that other households also contribute their fair share.

4.2. Experimental design

We sent letters to a sample of Dallas County homeowners and randomized some of the information contained in the letter. See a sample of the first page of the letter in Fig. 3. We designed the letters to include attributes that would signal legitimacy to recipients. For example, the letters were signed by a professor from The University of Texas at Dallas and the letters and envelopes included the official logo of The University of Texas at Dallas, a well-known institution in Dallas County.[14] We tailored the letters to each recipient, using their names and addresses. The letters also included contact information.[15]

All of our letters include a table in the middle of the first page, as shown in Fig. 3. In Fig. 3, this table is highlighted inside a red box with dashed lines (the box is for explanatory purposes and was not included in the actual letters sent to subjects). In all letters, the table includes the new information released by the DCAD on the household's own proposed value and estimated taxes for 2020 (i.e., the information in the second column of the table). The key randomization is whether the table also includes a third column with information about the proposed value and estimated taxes for the average household in Dallas County.

More precisely, we randomized (with a $\frac{2}{3}$ probability) whether the table includes a column that shows these figures for the average home in Dallas. This column allows recipients to compare whether their households' tax rates are above or below the average Dallas home and by how much. Randomly including this additional column is meant to provide a shock to households' perceptions of the tax rate that other households pay. We also cross-randomized an additional row that makes the tax rates explicit. We include this minor treatment arm because prior literature suggests that some people struggle with basic numerical calculations (e.g., Kahneman and Tversky (1979), Lipkus et al. (2001), Peters et al. (2006)) or may not be inclined to perform them during experiments (e.g., Kahneman et al. (1990), Frederick (2005), Chetty et al. (2009)). Since we cross-randomized two aspects of the tables, in total there were four possible types of tables, which are summarized in Fig. 4.

We created the letters as soon as administrative data became available (on May 16th, 2020), mailed them on May 20th, and started receiving responses to the Field Survey on May 21st. The post office scanned more than 90% of our letters by Friday, May 22nd, 2020, indicating that they reached the last mile before delivery by then. Historical data show that most homeowners file their protests near the deadline. As a result, there was ample time between receipt of the letter and the protest deadline for the letter to influence most recipients' decisions to protest.[16]

Our pre-registered field experiment included two main treatment arms. The first, described in this paper, focused on fairness. The second, detailed in a companion paper Nathan et al. (2025), was designed to measure hassle costs. In brief, one group received no letter, a second group received a letter with basic instructions on how to protest, and a third group received the letter plus an *extra aid* message.[17]

[14] See Appendix B for a sample letter and Appendix C for a sample envelope.

[15] Appendix A.2.1 presents more details about the letter, including the contact information. We included a URL to a website that provided help on how to protest (Nathan et al., 2025).

[16] Figure A.10 in Nathan et al. (2025) shows the timing of the survey responses, of the visits to the study's website, and of the protests.

[17] In our RCT pre-registration, we included several analyses—such as heterogeneity tests—related to the treatment arm examined in the companion paper Nathan et al. (2025). We conducted some of these analyses for the treatment arm studied in this paper, but the results are underpowered and therefore inconclusive.

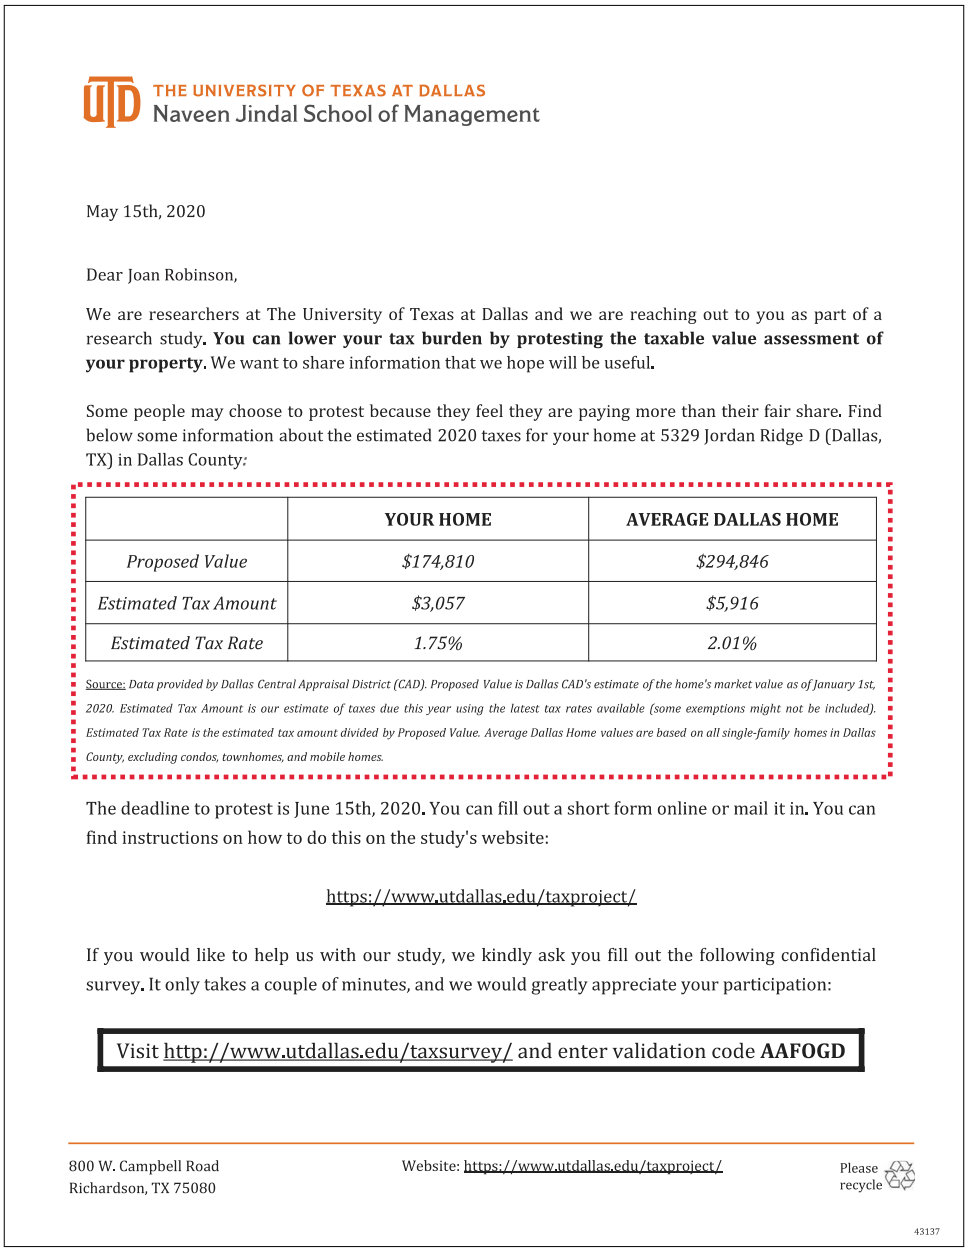

THE UNIVERSITY OF TEXAS AT DALLAS
Naveen Jindal School of Management

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

Fig. 3. First Page of the Letter from the Field Experiment
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 Fig. 4. The table appears inside a red frame with dashed lines (this frame was added to this figure for emphasis but does not appear in the actual letters).

4.3. Econometric specification

When our letters inform subjects that the average tax rate in Dallas County is 2.01%, subjects may update their beliefs upward or downward depending on whether their prior beliefs about the average tax rate were higher or lower than 2.01%. Given that the field experiment does not observe subjects' prior beliefs about the average tax rate, we leverage the systematic misperceptions documented in the Perceptions Survey. Specifically, we use the fact that homeowners who pay below-average tax rates tend to underestimate the average tax rate, while those who pay above-average tax rates tend to overestimate it.

The following econometric specification formalizes this empirical strategy. Let Y_i^{post} denote the outcome of interest. For example, this outcome could be an indicator variable that takes the value 100 if the household protested directly in 2020 (the post-treatment

a. Average: No, Rate: No

	YOUR HOME
Proposed Value	*$174,810*
Estimated Tax Amount	*$3,057*

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).*

b. Average: Yes, Rate: No

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

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). Average Dallas Home values are based on all single-family homes in Dallas County, excluding condos, townhomes, and mobile homes.*

c. Average: No, Rate: Yes

	YOUR HOME
Proposed Value	*$174,810*
Estimated Tax Amount	*$3,057*
Estimated Tax Rate	*1.75%*

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.*

d. Average: Yes, Rate: Yes

	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.*

Fig. 4. Sample Information Tables
Notes: Each panel corresponds to the hypothetical table that a given subject household would receive if assigned to the specified treatment group. The table was placed in the middle of the first page of the letter, within the red frame with dashed lines in Fig. 3 (the red frame was not included in the actual letters—see Appendix B for a sample of the letter without the red boxes added). Appendix A.2.1 provides details on how the values shown in each table were calculated.

period). Let τ_i represent the household's own tax rate, and let $\bar{\tau}^{feed}$ denote the *actual* average tax rate in the county. Let D_i be an indicator variable equal to 1 if the information on the average tax rate was shown to the subject in the experiment. The regression of interest is as follows:

$$Y_i^{post} = \nu_0 + \nu_1 \cdot D_i \cdot (\bar{\tau}^{feed} - \tau_i) + \nu_2 \cdot (\bar{\tau}^{feed} - \tau_i) + \nu_3 \cdot D_i + X_i^{pre} \nu_X + \varepsilon_i \quad (2)$$

The variable X_i^{pre} corresponds to the vector of additional control variables. Unless stated otherwise, we use the following set of controls in all regressions: the proposed value, the growth in the proposed value relative to the previous year, the number of years since the household's last property tax appeal, and indicator variables for homestead status, multiple owners, school districts, special districts, the extra aid message treatment (see Nathan et al. (2025)) and for whether the household appealed in each year from 2015 to 2020. In principle, an experiment should ensure balance in control variables. However, including controls can help address concerns that, by chance, the treatment and control groups may be unbalanced along some dimensions. A further reason why experimental studies often include control variables is that they can increase statistical power by reducing the variance of the error term (McKenzie, 2012). The coefficient ν_2 captures the relationship between the outcome and the potential information shock (i.e., $\bar{\tau}^{feed} - \tau_i$) when the information is not disclosed. The key coefficient, ν_1, measures the effect of the information shock—that is, how much stronger the relationship becomes when the feedback is disclosed.

4.4. Subject pool

Our subject pool of interest consists of a sample of 50,394 subjects to whom we sent a letter as part of the mailing intervention.[18] The main outcome of interest in the field experiment is whether the household filed a tax protest in 2020. Among households in the control group (i.e., those who received a letter that did not include information on the average household in Dallas County), 11.08% filed a direct protest, and an additional 5.85% protested using an agent.

Column (1) of Table 1 presents average pre-treatment characteristics of the subject pool. The average subject owns a home worth $343,162 and pays an effective tax rate of 2.10%. Columns (2) and (3) show that, consistent with successful random assignment, pre-treatment characteristics are largely balanced between the treatment and control groups.[19] Lastly, although the subject pool is not perfectly representative of the universe of 423,607 single-family homes in Dallas County, it is close to being so (see Appendix A.2.1 for details).

[18] See Appendix A.2.1 for the sample selection criteria. We sent letters to 50,983 households but removed 589 households from the sample because they had already protested by the time the letters were sent. This resulted in a final sample of 50,394 households.

[19] The difference is statistically significant for one of the pre-treatment variables: the home value. However, while the difference is statistically significant due to the large sample size, it is small in magnitude. Moreover, we report 10 tests in column (4)—by chance, one of them is expected be statistically significant at the 10% level.

Table 1
Randomization balance test: Field experiment.

	Full sample				Field survey subsample			
	(1) All	(2) Treatment	(3) Control	(4) P-value	(5) All	(6) Treatment	(7) Control	(8) P-value
Home Value ($1,000s)	343.162	345.235	339.001	0.034	394.983	395.021	394.916	0.994
	(1.406)	(1.743)	(2.371)		(6.987)	(8.548)	(12.092)	
Property Tax Rate (%)	2.103	2.104	2.100	0.364	2.107	2.125	2.075	0.014
	(0.002)	(0.003)	(0.004)		(0.010)	(0.012)	(0.016)	
Owner-Protest in 2019 (%)	5.854	5.934	5.692	0.273	10.328	10.370	10.253	0.936
	(0.105)	(0.129)	(0.179)		(0.701)	(0.875)	(1.170)	
Agent-Protest in 2019 (%)	4.598	4.543	4.708	0.407	5.667	5.926	5.201	0.507
	(0.093)	(0.114)	(0.164)		(0.532)	(0.678)	(0.857)	
2020 Homestead Exemption (%)	83.760	83.779	83.722	0.871	92.797	92.346	93.611	0.297
	(0.164)	(0.201)	(0.285)		(0.595)	(0.763)	(0.943)	
Number of Bedrooms	3.304	3.303	3.306	0.554	3.410	3.396	3.435	0.197
	(0.003)	(0.003)	(0.005)		(0.015)	(0.018)	(0.025)	
White (%)	44.275	44.412	44.000	0.380	56.568	57.449	54.978	0.301
	(0.221)	(0.271)	(0.383)		(1.141)	(1.419)	(1.919)	
Hispanic (%)	27.321	27.290	27.382	0.827	15.307	15.556	14.859	0.686
	(0.199)	(0.243)	(0.344)		(0.829)	(1.040)	(1.372)	
Black (%)	18.685	18.579	18.897	0.389	17.691	16.626	19.614	0.110
	(0.174)	(0.212)	(0.302)		(0.878)	(1.069)	(1.532)	
Asian (%)	9.719	9.719	9.720	0.997	10.434	10.370	10.550	0.903
	(0.132)	(0.162)	(0.229)		(0.704)	(0.875)	(1.185)	
Observations	50,394	33,635	16,759		1,888	1,215	673	

Notes: Average pre-treatment (i.e., before the start of letter delivery) characteristics of subjects in the field experiment, with standard errors in parentheses. Column (1) corresponds to the entire field experiment sample, column (2) contains subjects in the field experiment that received the average rate treatment, and column (3) contains subjects in the field experiment that did not receive the information on the average rate. Column (4) reports the *p*-value of the test of equal means across the two groups shown in columns (2) and (3). Column (5) corresponds to the entire Field Survey sample. Column (6) contains respondents to the Field Survey who received the average rate treatment, and column (7) contains respondents of the Field Survey that did not receive the information on the average rate. Column (8) reports the *p*-value of the test of equal means across the two groups shown in columns (6) and (7). *Home Value* is the proposed assessment value; *Property Tax Rate* is the ratio of the property tax amount over the home value; *Owner-Protest in 2019* and *Agent-Protest in 2019* indicate whether the subject protested directly or through an agent in 2019, respectively; *2020 Homestead Exemption* indicates an effective homestead exemption. *Number of Bedrooms* is the number of bedrooms in the respondent's home. *White*, *Hispanic*, *Black*, and *Asian* are the shares of homeowners by each imputed race.

4.5. Field Survey

The letters included a URL to an online survey, accessible only by entering a unique five-letter code. This code ensured that only recipients could respond and allowed us to match survey data to administrative records. Survey responses served as a proxy for households who closely engaged with the letter. In addition, the survey included a few questions that could be used as outcome variables.

A sample of the survey instrument is provided in Appendix G. One question asked respondents about the likelihood that they would protest in 2020, using a scale from 1 (very likely) to 4 (very unlikely). This question was designed to capture the "immediate" effect of the information in the letter on protest intentions, shortly after the information was received. Another question asked whether the respondent felt that his or her own taxes were fair, using a 1–10 scale: "Relative to the other households in the county, do you think your household pays a fair amount in property taxes?" This question was intended to test the fairness mechanism. The hypothesis is that someone who learns the average tax rate is higher than previously believed is more likely to view his or her own taxes as fair.

We received responses to the Field Survey from 1,888 households, implying a response rate of 3.7%. This is consistent with rates found in other studies using similar delivery methods (i.e., survey links provided via letter). For example, the response rate was 3.6% in Giaccobasso et al. (2025), and the meta-study by Sinclair et al. (2012) reports an average response rate of 4.7%. The median respondent took approximately 6.3 min to complete the survey.

When examining the results for the subsample of Field Survey respondents, it is important to note that they do not constitute a random sample. On the one hand, while differences in household characteristics—such as home value, number of bedrooms, and tax rate—between respondents and non-respondents are statistically significant, they are small in magnitude (results reported in Appendix A.2.4). However, there is a substantial difference in the share of subjects who protested in the post-treatment year (2020): 50.52% among survey respondents versus 11.08% in the overall sample who received letters. This difference has a natural interpretation: subjects who paid the most attention to our letter were likely those on the margin of deciding whether to protest in

Table 2
Results from the field experiment: Effects of the information about the average tax rate in the county.

	Full sample			Field survey		
	(1) P^d_{2020}	(2) P^d_{2020}	(3) P^d_{2020}	(4) P^d_{2020}	(5) I^d_{2020}	(6) F^d_{2020}
Panel A. With Controls						
Information Shock ($\overline{\tau}$)	−0.937*	−2.131*	−0.528	−12.566**	−11.919***	0.459*
	(0.509)	(1.224)	(0.524)	(5.424)	(4.495)	(0.244)
Mean Outcome (Control)	11.08	15.67	6.52	50.52	82.86	4.02
Std. Dev. Outcome (Control)	31.39	36.35	24.69	50.03	37.72	2.13
Observations	50,394	25,197	25,197	1,888	1,867	1,888
Panel B. Without Controls						
Information Shock ($\overline{\tau}$)	−0.835	−1.790	−0.416	−12.095**	−11.974***	0.499**
	(0.524)	(1.260)	(0.529)	(5.636)	(4.586)	(0.243)
Mean Outcome (Control)	11.08	15.67	6.52	50.52	82.86	4.02
Std. Dev. Outcome (Control)	31.39	36.35	24.69	50.03	37.72	2.13
Observations	50,394	25,197	25,197	1,888	1,867	1,888

Notes: Robust standard errors in parentheses. All columns present results from Eq. (2) in Section 4.3, with panel (a) presenting results with control variables and panel (b) showing results without controls. The variable *Information Shock* ($\overline{\tau}$) corresponds to the information shock term ($D_i \cdot (\overline{\tau}^{feed} - \tau_i)$). Column (1) reports results for subjects in the field experiment who received a letter. Column (2) shows results for subjects who were predicted by a probit model to have an above-median probability of responding to the Field Survey contained in the mailed letters, while column (3) shows the same for subjects with below-median probability of responding. Columns (4) through (6) report results for subjects who received a letter in the field experiment and responded to the Field Survey. The dependent variables are defined as follows: P^d_{2020} is an indicator variable that takes the value 100 if the owner protested directly in 2020 and 0 otherwise; I_{2020} is an indicator variable that takes the value 100 if the household reported to be either likely or very likely to protest in 2020 and zero otherwise; F_{2020} corresponds to a question about whether the taxes of the respondent are fair relative to the taxes of everyone else, on a 1–10 scale. The regressions in panel (a) include the following controls: the 2020 proposed value in levels and its annual growth, dummies for multiple owners, school and special districts, number of years since the household's last protest, a dummy for homestead status, a dummy indicating if the household received the extra aid message, and, for each previous year since 2015, a dummy indicating if the household protested in that year and the outcome of the protest (if any) as a percent-reduction in the market value (i.e., the protest history). Control variables for the protest history depend on the year in which the dependent variable is measured. For instance, if the outcome corresponds to direct protests in 2018, the protest history controls include protests in 2015, 2016, and 2017. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively, based on two-sided tests.

2020.[20] Given that survey participation is self-selected, it is also important to examine balance within the survey subsample. Columns (5) through (8) of Table 1 present these results. Consistent with successful random assignment, pre-treatment characteristics are balanced between the treatment and control groups.[21]

5. Results from the field experiment

5.1. Main results

Table 2 presents the results of the field experiment based on estimates from Eq. (2). Panel (a) reports results with controls, while panel (b) shows results without any controls. Column (1) uses the full sample of 50,394 subjects who were randomly selected to receive a letter. The dependent variable, drawn from administrative data, takes the value 100 if the household protested directly in 2020 and 0 otherwise. The coefficient on the *Information Shock* ($D_i \cdot (\overline{\tau}^{feed} - \tau_i)$) in panel (a) indicates that learning the average tax rate ($\overline{\tau}$) is 0.1 pp higher causes a 0.093 pp reduction in the probability of protesting. This effect is statistically significant (p-value = 0.066).

A common concern in correspondence experiments is that households may not have opened, or paid attention to, the letters sent to them. Evidence from a companion paper, Nathan et al. (2025), suggests otherwise. Specifically, individuals who received a letter with basic information on how to file a protest were 1.792 pp more likely to file an appeal than those who did not receive a letter, and appeal rates were even higher among those who received a letter with additional assistance on protesting.[22] This provides direct evidence that a non-negligible share of households in this field experiment read the letters. However, this evidence does not guarantee that *every* household paid attention.

The *Information Shock* coefficient reflects an intention-to-treat effect. There are multiple potential sources of non-compliance—the first being that many households may not have read our piece of unsolicited mail. Imperfect attention is a well-documented fact in mailing interventions more broadly (e.g., Perez-Truglia and Cruces (2017), Gerber et al. (2020), Bottan and Perez-Truglia (2025), Nathan et al. (2025)). For instance, Perez-Truglia and Cruces (2017) estimates that only 21.5% of recipients actually absorbed the

[20] Our letter makes clear from the outset that it concerns tax protesting, so individuals considering a protest were more likely to continue reading. Moreover, those who paid attention may have found the information helpful for submitting a protest.

[21] There is a statistically significant but small difference in one of the ten pre-treatment variables: the property tax rate.

[22] See column (1) of Table 2 in Nathan et al. (2025).

relevant information contained in the letter sent to them. As a result, these studies often find that, due to substantial non-compliance, treatment-on-the-treated effects are significantly larger than the corresponding intention-to-treat effects.[23]

We use two approaches to account for this form of non-compliance. Both rely on the fact that, for subjects who responded to the Field Survey, we have confirmation that they paid attention to the letter—otherwise, they could not have accessed the survey link and code, which were provided only in the personalized letter. Indeed, assuming subjects read the letter from top to bottom, they likely noticed the information about the average tax rate, as the survey link appeared directly below the table containing the information shock.

The first strategy focuses on identifying the types of individuals most likely to respond to the Field Survey. We estimate a simple probit model to compute the ex-ante probability that a household responds, using a range of pre-treatment characteristics such as the household's protest history, proposed value, tax rate, and homestead status (see Appendix A.2.6 for details). The model has decent predictive power, with an out-of-sample area under the receiver operating characteristic curve (AUC) of 0.66.

In the next step, we split the sample based on whether the ex-ante probability of responding is above or below the median and then estimate the same regression as in column (1) of Table 2 separately for each half of the sample. Table 2 reports these results: column (2) of panel (a) corresponds to subjects with an above-median response probability, and column (3) to those with a below-median probability. Subjects in column (2) are more than twice as likely to respond to the survey as those in column (3) (5.6% vs. 1.9%). They are also more likely to protest: the average protest rate is 15.6% in column (2) versus 6.5% in column (3). This is consistent with the idea that individuals who were more interested in the letter were also more likely to be contemplating protesting. Comparing columns (2) and (3) suggests that, as expected, the effects of the information shock are stronger among individuals more likely to pay attention. Specifically, the coefficient is −2.131 (p-value = 0.082) in column (2) versus −0.528 (p-value = 0.313) in column (3). However, this difference should be interpreted with a grain of salt, as the estimates are imprecise and the difference between the two coefficients is statistically insignificant (p-value = 0.228). Panel (b) presents the same results without controls. The coefficients in columns (1) through (3) are similar in magnitude with controls (panel (a)) and without controls (panel (b)). However, the findings in panel (b) are weaker, as the coefficients in columns (1) and (2) become borderline insignificant when controls are removed (p-values = 0.111 and 0.155, respectively).

As a second (and more direct) strategy, we focus on the subjects who responded to the Field Survey. However, since the decision to respond occurs in the post-treatment period, conditioning on this post-treatment variable may contaminate the experimental variation. For this reason, when restricting the analysis to survey respondents, it is important to conduct appropriate falsification tests, such as an event-study analysis.

Column (4) of Table 2 uses the same specification as column (1), except that it is restricted to the subsample of 1,888 households who responded to the Field Survey. The coefficient in column (4) of panel (a) is negative (−12.566) and statistically significant (p-value = 0.021). It is also much larger in magnitude than the corresponding coefficient in column (1). This difference is partly mechanical: since the baseline protest rate is much higher among survey respondents (50.52% in column (4) vs. 11.08% in column (1)), larger effect sizes are expected. However, this mechanical difference does not fully explain the result. While the average outcome is 4.5 times larger (50.52 vs. 11.08), the estimated effect is 13.4 times larger (−12.566 vs. −0.937). We interpret this as evidence that the effects are stronger in column (4) largely because those households paid more attention to the information in the letter. Column (4) of panel (b) shows that the estimated coefficient and overall conclusions remain similar even when control variables are excluded.

The results from column (4) of panel (a) suggest that learning the average tax rate is 0.1 pp higher leads to a 1.256 pp decrease in the probability of protest. The most appropriate benchmark for comparison is the average protest rate of 11.08% among households in the control group (i.e., those who received a letter that did not include a column for the average household). Finding that the average tax rate is 0.1 pp higher reduces subjects' likelihood of protesting by approximately 11.3% ($= \frac{1.256}{11.08}$) relative to that benchmark.[24]

5.2. Robustness checks

A fundamental concern in experimental studies is the possibility of spurious effects. To address this, and following previous work (Nathan et al., 2025; Giaccobasso et al., 2025), we conduct an event-study analysis. Since the information was provided in 2020, it cannot causally influence the decision to file an appeal in any year before 2020. Therefore, if effects appear not only in post-treatment years but also in pre-treatment years, they should be considered spurious.

The results are shown in Fig. 5(a). The rightmost coefficient in the figure shows the effect of the information shock on the probability of protesting in 2020 and corresponds to the coefficient reported in column (1), panel (a) of Table 2. The remaining coefficients are estimated using the same regression specification but with protests in each of the pre-treatment years (2015–2019) as the dependent variable.[25] The coefficients for each pre-treatment year are close to zero, statistically insignificant, and precisely estimated. This evidence indicates that the effect is not spurious.

[23] In Appendix A.2.5, we provide additional details about non-compliance in mailing experiments.

[24] An alternative benchmark would be the mean protest rate of 8.67% among households who were randomly selected not to receive a letter (from column (1) of Table 2 in Nathan et al. (2025)). Relative to this benchmark, the effect would be even larger: 14.5% ($= \frac{1.256}{8.67}$).

[25] The regressions in this figure use rolling periods of protest history as controls. For example, when the dependent variable is protest activity in 2019, we control for indicator variables representing protest history from 2015–2018. We include protest history starting in 2015, as the DCAD website provided data going back to that year at the time of the field experiment.

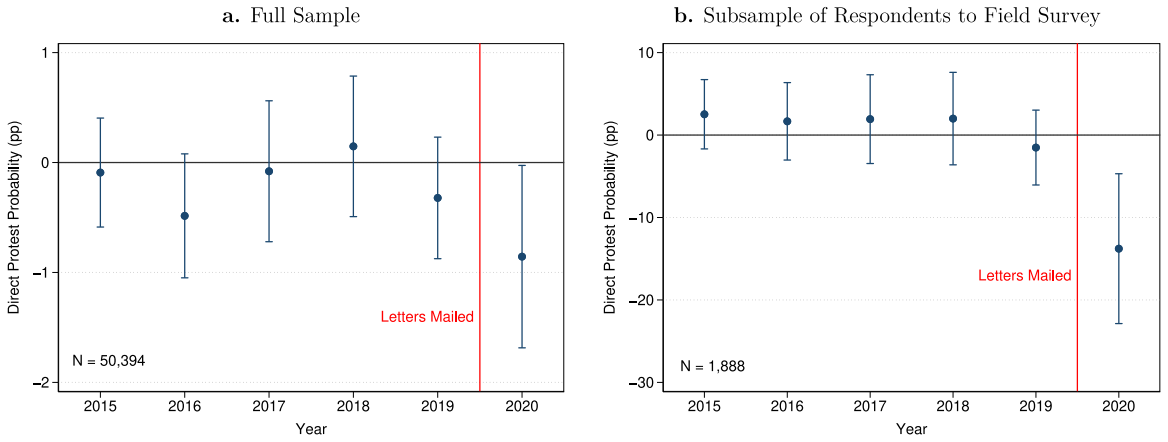

Fig. 5. Results from the Field Experiment: Event-Study Falsification Tests

Notes: Point estimates with 90% confidence intervals in brackets, based on robust standard errors. The point estimates are computed in the same way within both panels: the point estimates within each panel only change the focal year. Panel (a): The blue dots represent the coefficient on the information shock ($D_i \cdot (\overline{\tau}^{feed} - \tau_i)$) based on Eq. (2) from Section 4.3, with direct protests as the dependent variable for the sample of 50,394 subjects in the field experiment who received a letter. Panel (b): same as panel (a) except that it is based on the subsample of 1,888 subjects who responded to the Field Survey. The regressions in this figure include the following controls: the 2020 proposed value in levels and its annual growth, dummies for multiple owners, school and special districts, number of years since the household's last protest, a dummy for homestead status, a dummy indicating if the household received the extra aid message, and, for each previous year since 2015, a dummy indicating if the household protested in that year and the outcome of the protest (if any) as a percent-reduction in the market value (i.e., the protest history). Control variables for the protest history depend on the year in which the dependent variable is measured. For instance, if the outcome corresponds to direct protests in 2018, the protest history controls include protests in 2015, 2016, and 2017.

Regarding the results from column (4) of panel (a) from Table 2 for the Field Survey sample, a potential concern is that, despite random assignment, selection into the survey is endogenous and could introduce selection bias. To address this, Fig. 5(b) presents the event-study analysis for the Field Survey subsample. Specifically, we estimate the same regression as in column (4) of panel (a), but use as dependent variables indicators for whether the respondent protested in each of the pre-treatment years. Reassuringly, the estimated effects for the pre-treatment years are close to zero, statistically insignificant, and precisely estimated.

With the subsample of respondents to the Field Survey, we can estimate the effects of the information shock on the survey outcomes. Column (5) of panel (a) from Table 2 is identical to column (4) of panel (a) except that the dependent variable takes the value 100 if the household reports that it is likely or very likely to protest in 2020, and 0 if the household reports that it is unlikely or very unlikely to protest. There are reasons why the effects on stated intentions may differ from the effects on actual protests. For example, some individuals may care about the information provided but forget about it by the time they must decide whether to protest. In addition, the information may increase the desire to file a protest, but the individual may later fail to act as planned due to the difficulty of the process or other filing frictions (Nathan et al., 2025). In fact, we do not expect the effects on intended and actual protests to be mechanically the same, as the correlation between these two outcomes is statistically significant (p-value < 0.001) but far from perfect (correlation coefficient of 0.398). We find that the effects on the intention to protest closely mirror the effects on actual protests: the coefficient from column (5) of panel (a) is negative (−11.919) and highly statistically significant (p-value = 0.008). This coefficient is close in magnitude to and statistically indistinguishable from the corresponding coefficient of −12.566 from column (4) of panel (a). This result is essentially identical in column (5) of panel (b) when omitting controls.

In Appendix A.2.6, we present further robustness checks and results. We show that the effects of the information shock on protest choices were consequential for households' assessed values and tax amounts. The effects also persist when extending the horizon to include the following two years, indicating that the treatment did not simply shift the timing of protests that would have occurred anyway. We do not find significant differences between letters that made the tax rate explicit in a separate row of the table and those that did not. We also examine effects on protests filed through agents, and the evidence suggests there were likely not any effects on this outcome.[26] We show that the findings remain robust when restricting the sample to subjects who indicated on the Field Survey that they consider it fair for everyone to pay exactly the same tax rate. Finally, we show that a linear specification provides a reasonable approximation and that the results are not driven by outliers.

5.3. *The fairness channel*

The evidence above suggests that information about the average tax rate significantly influenced the decision to protest. Next, we use data from the Field Survey to provide suggestive evidence in favor of the fairness channel.

[26] In the full sample, the estimated effects are close to zero and statistically insignificant. While some significant effects appear in the Field Survey subsample, the event-study analysis indicates these are most likely spurious.

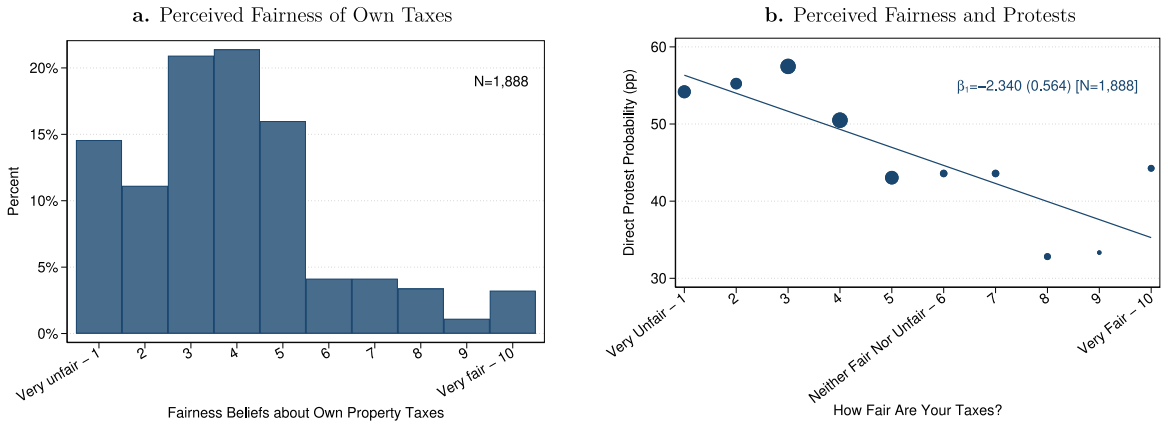

Fig. 6. Descriptive Results from the Field Survey

Notes: Non-experimental evidence of the fairness channel based on responses from 1,888 subjects in the Field Survey. Panel (a) features responses from subjects who answered the question, "Relative to the other households in the county, do you think your household pays a fair amount in property taxes?" Subjects could answer on a scale from 1 to 10, with 1 indicating "Very unfair" and 10 indicating "Very fair." Panel (b) presents a binned scatterplot of the association between direct protest rates (observed in the administrative records) and households' perceived fairness of their own property taxes (from a question included in the Field Survey). The size of each circle is proportional to the number of subjects that provided the response indicated. The line corresponds to the linear fit and is shown with the corresponding slope and robust standard error (in parentheses).

The survey included a question to measure whether households felt their own taxes, relative to the county average, were fair. Before measuring the treatment effects on this outcome, we present some basic descriptive analysis. Fig. 6(a) shows that there is substantial variation across households in their perceptions of fairness.[27] Fig. 6(b) examines whether perceived fairness correlates with the probability of filing a tax protest. There is a strong, negative, and statistically significant (p-value < 0.001) association between perceived fairness (x-axis) and the probability of filing a tax appeal (y-axis). An increase in fairness perceptions from very unfair (1) to very fair (10) is associated with a substantial 23.40 pp decrease in the likelihood of protesting.

According to the fairness channel, we would expect that an information shock that increases the perceived average tax rate would have a positive effect on perceived fairness. That is, if individuals learn that others are contributing their fair share, they should feel more obligated to contribute their own. The relevant results are presented in column (6) of Table 2, which is identical to column (4) except that the dependent variable is the household's perceived fairness of its own taxes. As expected, the coefficient on *Information Shock* is positive and statistically significant: 0.459 in panel (a) (p-value = 0.060) and 0.499 in panel (b) (p-value = 0.040). The effect on perceived fairness is comparable in magnitude to the effect on the probability of protest. The results from column (6) of panel (a) indicate that an information shock of 0.1 pp causes a change in perceived fairness equivalent to 2.15% ($= \frac{0.1 \cdot (-0.459)}{2.13}$) of a standard deviation. In comparison, the results from column (4) of panel (a) indicate that an information shock of 0.1 pp leads to an increase in the probability of protest equivalent to 2.48% ($= \frac{0.1 \cdot (-12.566)}{50.52}$) of a standard deviation.[28]

Together, the evidence in this section suggests that the effect of the information shock on protests operates through the fairness channel. An alternative mechanism—which we refer to as self-interest—posits that subjects may have reacted to information about the average tax rate because they inferred whether their own protests were likely to succeed. In Appendix A.2.7, we present evidence against this alternative mechanism.

5.4. Magnitude of the effects

One challenge in assessing the magnitude of fairness concerns is the need to account for two forms of non-compliance. The first is that some recipients may not have read the letter. As previously explained, we address this by focusing on the results for the subsample of respondents to the Field Survey (column (4) of Table 2, panel (a)). However, a second form of non-compliance remains: even if recipients read the letter, they may not fully incorporate the information into their beliefs. To correct for this, we use results on the strength of belief updating from the Perceptions Survey. Specifically, each additional 1 pp in the information shock increased the perceived average tax rate by 0.455 pp (i.e., the difference in slopes between treatment and control groups in Fig. 2(b)), implying a scaling factor of 2.19 ($= \frac{1}{0.455}$). Scaling up the coefficient on *Information Shock* from column (4) of panel (a) from Table 2 implies that increasing recipients' perception of the average county tax rate by 0.1 pp would decrease their protest probability by 2.75 pp ($= 0.1 \cdot 12.566 \cdot 2.19$).

[27] The distribution is skewed, with a majority of households believing that their taxes are unfair. For instance, in the control group, the average fairness score is 4.02, which is closer to the unfair end of the scale.

[28] In Appendix A.2.6, we use a binned scatterplot to show that the results do not appear to be driven by outliers or non-linearities.

To further illustrate the magnitude of the effects, we express the willingness to pay for fairness in monetary terms. These back-of-the-envelope calculations combine the results from this paper with those from Nathan et al. (2025). Specifically, Nathan et al. (2025) leverages a kink in the marginal benefits from protesting, created by the Texas Property Code's homestead cap, which limits annual increases in assessed values to 10%. Nathan et al. (2025) shows that a reduction in expected tax savings of $100 leads to a 2.14 pp decrease in the probability of protest.[29] Combining this with our finding that a 0.1 pp increase in the perceived average tax rate decreases the protest probability by 2.75 pp, we estimate that a 0.1 pp increase in the perceived average tax rate (equivalent to an increase of $295 in the average tax amount) raises the willingness to pay taxes by about $128 ($= \frac{2.75 \cdot 100}{2.14}$).[30]

The calculations above indicate that for each additional dollar in taxes contributed by the average household, a homeowner is willing to pay an additional $0.43 ($= \frac{128}{295}$) in their own taxes. As a benchmark, we can compare this estimate with findings from laboratory experiments. Specifically, variants of the public goods game show that individuals are more willing to contribute to a public good when they believe others are contributing as well (Fehr and Schmidt, 1999; Fischbacher et al., 2001; Gachter, 2007). Ex-ante, it is not obvious that fairness considerations would carry significant weight in the real-world context of tax compliance. For instance, the stakes are orders of magnitude higher for property taxes than in laboratory games, and taxpayers may hold strong views about the role of government that are largely absent in lab settings (Huet-Vaughn et al., 2019). Yet despite these differences, we observe a striking similarity between our field experiment results and lab-based estimates. For example, Figure 1 in Fischbacher et al. (2001)—one of the earliest studies on conditional cooperation—shows that for each additional dollar contributed by others, the average subject increases their own contribution by $0.45. This $0.45 estimate is remarkably close to our corresponding estimate of $0.43 from our field experiment.

6. Preferences Survey

In our field experiment, we informed households about the tax rate paid by the average household. Some participants learned they paid more or less than the average and adjusted their fairness perceptions and willingness to pay taxes accordingly. However, our information treatment did not explain *why* their tax rate differed from the average. In this section, we use a complementary survey experiment to explore whether—and how—the reasons behind these differences might influence the fairness effects.

6.1. Research design and results

Exemptions and caps—which account for most of the variation in property tax rates—may arguably be rooted in principles of fairness. For instance, certain households, such as those with disabilities, may receive exemptions precisely because doing so is considered fair. This has implications for interpreting the findings from the field experiment: while taxpayers may perceive it as unfair that a household pays a lower tax rate when no explanation is provided, their views might be less negative if they knew the reduction was due to a disability, for example.

To explore these fairness ideals further, we conducted a complementary survey experiment. We recruited participants on Prolific, restricting participation to U.S. homeowners over the age of 21. We collected responses from 528 participants, with 99% passing the attention check. The survey included two main modules. We randomized both the order in which the modules were presented and the order of the scenarios within each module. A sample of the survey instrument is provided in Appendix H. Appendix A.3 offers additional details on the survey's design and implementation, which are summarized below.

In the first module, we asked respondents to imagine that the government needed to collect a specific amount of annual property tax revenue from two households (Households A and B). Respondents were then asked to decide how much of the required tax revenue each household should pay. Using two sliders, they could set the tax rate for one household, while the rate for the other adjusted automatically and instantly to meet the revenue target, simplifying the task.

The results from this module are presented in Fig. 7(a). The histogram illustrates how respondents allocated tax rates between Households A and B, categorized by scenario. For example, the middle category (labeled "-0.25 to +0.25") represents the proportion of respondents who assigned approximately equal tax rates to both households.[31]

The left-most blue bar in each bin represents the control scenario, in which Households A and B have the same home value and no distinguishing characteristics. The vast majority of respondents (86.0%) chose to assign approximately the same tax rate to both households. The remaining bars correspond to four alternative scenarios. In the next scenario (labeled *Homestead*), the two households again have identical home values, but one is a primary residence while the other is not. In this case, the share of respondents who equalized tax rates drops from 86.0% in the control group to 34.9% in the homestead group. In other words, a substantial fraction still prefers to assign the same tax rate to both households, despite their difference.

The results follow a similar pattern across the other three scenarios. In the *Over Age 65* scenario, one household's owner is over 65, while the other is not. In the *Disabled* scenario, one household has an owner with a disability, whereas the other does not. In the *Poorer* scenario, the only difference lies in the home values, with one household owning a significantly lower-value home ($150,000

[29] This result is presented in Figure 1 of Nathan et al. (2025).

[30] The average home in Dallas County is worth $294,846 and pays $5,916 in property taxes. A change in the tax rate of 0.1 pp corresponds to a $295 change in the average tax amount ($= \$294{,}846 \cdot 0.001$).

[31] Most respondents in this category selected identical tax rates. For instance, in the control group (blue bar), of the 86.0% in the "-0.25 to +0.25" category, 97.1% chose rates within 0.05 pp of each other. Others likely intended to select the exact same rate but ended up with slight differences due to the difficulty of precisely adjusting the sliders.

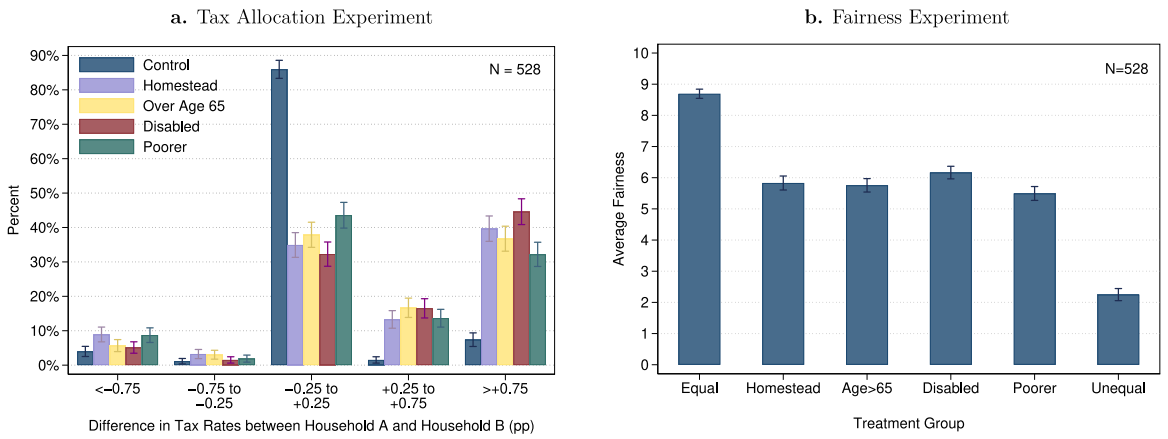

Fig. 7. Results from the Preferences Survey: Redistribution Preferences and Fairness

Notes: This figure shows responses to questions included in the Preferences Survey regarding the distribution of property tax burden. The survey was distributed to 528 homeowners in the U.S. over the age of 21. Panel (a) shows how respondents indicated that they would distribute a tax burden between Household A and Household B in five different scenarios. In the first four scenarios, subjects responded how they would distribute $18,000 between Household A and Household B, and in the fifth scenario subjects responded how they would distribute $12,000 between Households A and B. These five scenarios are summarized as follows. *Control*: Households A and B are identical. *Homestead*: Households A and B are identical, but the owner of Household A lives in their property as their primary residence and the owner of Household B does not live in their property as their primary residence. *Over age 65*: Households A and B are identical, but the owner of Household A is over age 65 and the owner of Household B is under age 65. *Disabled*: Households A and B are identical, but the owner of Household A is disabled and the owner of Household B is not disabled. *Poorer*: Household A's home is valued at $150,000, and Household B's home is valued at $450,000. The horizontal axis shows bins for the difference in tax rates between Household A and Household B in pp. The vertical axis presents the share of the responses for each bin. Panel (b) shows how respondents rated the fairness of different tax distribution scenarios between Households A and B on a scale from 0 (very unfair) to 10 (very fair). The horizontal axis corresponds to six scenarios which are summarized as follows. *Equal*: Households A and B are identical, each own a home worth $450,000, and both pay a tax rate of 2%. *Homestead*: Households A and B are identical and each own a home worth $450,000, but the owner of Household A (who lives in their property as their primary residence) pays a tax rate of 1.5% and the owner of Household B (who rents out their property) pays a tax rate of 2%. *Age* > 65: Households A and B are identical and each own a home worth $450,000, but the owner of Household A (who is over age 65) pays a tax rate of 1.5% and the owner of Household B (who is under age 65) pays a tax rate of 2%. *Disabled*: Households A and B are identical and each own a home worth $450,000, but the owner of Household A (who is disabled) pays a tax rate of 1.5% and the owner of Household B (who is not disabled) pays a tax rate of 2%. *Poorer*: Household A owns a home valued at $150,000 and pays a tax rate of 1.5%, while Household B owns a home valued at $450,000 and pays a tax rate of 2%. *Unequal*: Households A and B are identical and each own a home worth $450,000, but Household A pays a tax rate of 1.5% and Household B pays a tax rate of 2%. The vertical axis shows the average fairness rating for each scenario. Point estimates with 90% confidence intervals in brackets in both panels. See Section 6.1 for details on the design and implementation of the Preferences Survey.

vs. $450,000). Across all these scenarios, the findings are consistent: while most respondents are willing to provide tax breaks to households with certain characteristics, a substantial share (between 32% and 44%) still prefers that all households pay the same tax rate, regardless of individual circumstances.

In the first module discussed above, we asked respondents to select tax rates freely, without mentioning the word *fairness*. The second module was designed to measure perceived fairness more directly. In this set of questions, we described two households (Households A and B) and the tax rates they face. Respondents were then asked to assess the fairness of the tax distribution on a scale from 0 to 10, where 0 represents "very unfair" and 10 represents "very fair." The results from this module are presented in Fig. 7(b), which reports the average fairness score across six different scenarios. Two of these serve as benchmarks, shown by the leftmost and rightmost bars in the figure. In one benchmark, labeled *Equal*, Households A and B are identical and pay the same tax rate. As expected, respondents view this distribution as very fair, assigning it an average score of 8.7 on the 0–10 scale. In the other benchmark, labeled *Unequal*, Households A and B are identical but face different tax rates—1.5% for one and 2.0% for the other. As expected, respondents perceive this scenario as unfair, with an average fairness rating of 2.3, placing it near the bottom of the scale.

In the remaining scenarios presented in Fig. 7(b), one household pays a lower tax rate than the other (1.5% vs. 2.0%), but the household with the lower rate has a special circumstance. In the *Homestead* scenario, the household uses its home as a primary residence; in *Age* > *65*, the household owner is over 65; in *Disabled*, the household has a member with a disability; and in *Poorer*, the household owns a significantly lower-value home ($150,000 vs. $450,000). Comparing the *Unequal* and *Equal* scenarios in Fig. 7(b) shows that, in the absence of contextual information, unequal tax rates reduce perceived fairness by about 6.4 points (= 8.7 − 2.3). In contrast, when the context of a homestead exemption is provided, the unequal tax rates reduce perceived fairness by only 2.9 points (= 8.7−5.8). In other words, the fairness penalty associated with unequal tax rates is roughly halved when a justifying context is included (2.9 vs. 6.4 points, or a 55% reduction). This pattern holds across the other scenarios in which alternative justifications are provided for why one household faces a lower tax rate—*Over Age 65*, *Disabled*, and *Poorer*.

This final result from the Preferences Survey has important implications for interpreting the findings from the field experiment. In the field experiment, participants were not given any explanation for why the average tax rate was higher or lower than their own. If all subjects correctly inferred the underlying reasons—such as homestead or disability exemptions—we would expect the results to be unchanged even if that context had been made explicit. However, if some participants failed to make this inference, the findings would likely have differed had the additional context been provided explicitly. To illustrate, consider the extreme case in which no participants in the field experiment inferred the reasons for the differences in tax rates. In that case, we can use the Preferences Survey to adjust the field experiment's results. In the field experiment, learning that the average tax rate was 0.1 pp lower than expected reduced perceived fairness by an average of 0.459 points. According to the Preferences Survey, if participants had known that this lower rate reflected targeted exemptions, the drop in perceived fairness would likely have been only about half as large. Consequently, we would expect the corresponding effect on tax compliance to also be approximately half as strong.

7. Conclusions

This paper provides novel causal evidence on the role of fairness perceptions in tax compliance through a natural field experiment on U.S. property taxes. We mailed letters to households to introduce exogenous variation in their perceptions of the average tax rate paid by other households in their county. We then examined how this information shock affected households' high-stakes decisions to file a tax appeal, measured through administrative records, as well as their perceptions of tax fairness, measured through surveys. We show that when households perceive the average tax rate in the county to be higher (lower), their sense of fairness increases (decreases). Additionally, their willingness to pay taxes rises (falls), as reflected by a lower (higher) probability of filing a tax appeal. We estimate that for each additional $1 contributed by the average household, a taxpayer is willing to pay an extra $0.43 in their own taxes. In the main experiment, participants were told the average tax rate but received no explanation for why it differed from their own. In a complementary survey experiment, we demonstrate that this context matters: when households are informed that lower tax rates stem from exemptions—such as those for disability or senior status—they are more accepting of the differences in tax rates. The findings from this complementary survey suggest that future research should further examine how contextual factors and information framing influence fairness perceptions and tax compliance.

To our knowledge, this is the first study to provide causal, revealed preference evidence on the role of fairness considerations in tax compliance in a natural context. A key factor in any empirical study is the generalizability of its findings. In the language of List (2020), our results can be considered a wave-1 insight—establishing initial causality and providing the first tests of theory. We hope that future research will further explore how these findings vary across different contexts. We conclude by discussing key factors to consider when extrapolating our results.

First, our study focused on a single county due to the logistical costs of implementing the mailing intervention across multiple counties. Since property tax systems function similarly in other U.S. counties, we believe the results from our field experiment should generalize to other counties. However, when applying this framework elsewhere, political and cultural factors may mediate the effects. For instance, if reciprocal fairness norms are stronger (weaker) in some societies, we would expect the effects observed in our field experiment to be stronger (weaker) as well.

Second, our study examines a single type of tax: property taxes. When considering fairness concerns in other tax systems, significant differences may arise. The Preferences Survey provides suggestive evidence on this issue. We find that, for property taxes, a substantial portion of respondents prefer all households to pay the same tax rate, regardless of differences in home values. This evidence aligns with the fact that property tax schedules are approximately proportional, not only in Dallas County but across the United States (Dray et al., 2023). To explore whether these fairness preferences extend to other types of taxes, we included a similar question regarding income taxes.[32] Consistent with previous research (e.g., Ballard-Rosa et al., 2017), we find greater support for progressive schedules in the case of income taxes. However, even for income taxes, a significant share of respondents still prefer a proportional tax schedule.[33] This evidence suggests that while the magnitude of fairness effects may vary across tax types, the underlying qualitative patterns may hold beyond property taxes.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jacceco.2025.101838.

[32] Results presented in Appendix A.3.3.

[33] While we cannot determine the exact cause of these differences, our best guess is that they stem from historical precedent: property tax schedules have traditionally been proportional, whereas income tax schedules have historically been progressive.

References

Allingham, M.G., Sandmo, A., 1972. Income Tax Evasion: A Theoretical Analysis. J. Public Econ. 1, 323–338.

Alm, J., Bloomquist, K.M., McKee, M., 2017. When you know your neighbour pays taxes: Information, peer effects and tax compliance. Fisc. Stud. 38 (4), 587–613.

Alm, J., Torgler, B., 2011. Do ethics matter? Tax compliance and morality. J. Business Ethics 101 (4), 635–651.

Ballard-Rosa, C., Martin, L., Scheve, K., 2017. The structure of American income tax policy preferences. J. Politics 79 (1), 1–16.

Baugh, B., Ben-David, I., Park, H., 2018. Can taxes shape an industry? Evidence from the implementation of the Amazon tax. J. Financ. 73 (4), 1819–1855.

Belnap, A., Thornock, J., Standridge, K., Williams, B., 2025. Social tax avoidance.

Belnap, A., Welsch, A., Williams, B., 2023. Remote tax authority. J. Account. Econ. 75 (2–3), 101570.

Bergolo, M., Ceni, R., Cruces, G., Giaccobasso, M., Perez-Truglia, R., 2018. Misperceptions about Tax Audits. AEA Pap. Proc. 108, 83–87.

Bottan, N.L., Perez-Truglia, R., 2025. Betting on the House: Subjective Expectations and Market Choices. Am. Econ. J.: Appl. Econ. 17 (1), 459–500.

Cavallo, A., Cruces, G., Perez-Truglia, R., 2017. Inflation expectations, learning, and supermarket prices: Evidence from survey experiments. Am. Econ. J.: Macroecon. 9 (3), 1–35.

Chetty, R., Looney, A., Kroft, K., 2009. Salience and Taxation: Theory and Evidence. Am. Econ. Rev. 99 (4), 1145–1177.

Chow, T., Fan, Z., Huang, L., Li, O.Z., Li, S., 2023. Reciprocity in corporate tax compliance – evidence from ozone pollution. J. Account. Res. 61 (5), 1425–1477.

Congressional Budget Office, 2020. The Budget and Economic Outlook: 2020 to 2030. Technical Report, Congressional Budget Office.

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

Cullen, J.B., Turner, N., Washington, E., 2021. Political alignment, attitudes toward government, and tax evasion. Am. Econ. J.: Econ. Policy 13 (3), 135–166.

Dallas County Elections Department, 2024. General Election 2024 - Dallas County Vote Totals. Accessed At: www.Dallascountyvotes.Org/Wp-Content/Uploads/2024/11/Upload-6.Pdf.

DCAD, 2020. Data Products - Dallas Central Appraisal District. Accessed At: www.Dallascad.Org/DataProducts.Aspx. (Last Accessed: December 2020).

Del Carpio, L., 2022. Are the neighbors cheating? Evidence from a social norm experiment on property taxes in Peru. Working Paper.

Dobay, N., Nicely, F., Sanderson, A., Sanderson, P., 2019. The Best (And Worst) of International Property Tax Administration. Technical Report, Council on State Taxation, pp. 1–104.

Dray, S., Landais, C., Stantcheva, S., 2023. Wealth and property taxation in the United States. NBER Working Paper No. 31080.

Fehr, E., Schmidt, K.M., 1999. A Theory of Fairness, Competition, and Cooperation. Q. J. Econ. 114 (3), 817–868.

Fehr, E., Schurtenberger, I., 2018. Normative foundations of human cooperation. Nat. Hum. Behav. 2 (7), 458–468.

Fischbacher, U., Gachter, S., Fehr, E., 2001. Are people conditionally cooperative? Evidence from a public goods experiment. Econom. Lett. 71 (3), 397–404.

Floyd, E., List, J.A., 2016. Using field experiments in accounting and finance. J. Account. Res. 54 (2), 437–475.

Frederick, S., 2005. Cognitive reflection and decision making. J. Econ. Perspect. 19 (4), 25–42.

Gachter, 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. MIT Press, Cambridge, MA, US.

Gerber, A., Hoffman, M., Morgan, J., Raymond, C., 2020. One in a million: Field experiments on perceived closeness of the election and voter turnout. Am. Econ. J.: Appl. Econ. 12 (3), 287–325.

Giaccobasso, M., Nathan, B., Perez-Truglia, R., Zentner, A., 2025. Where Do My Tax Dollars Go? Tax Morale Effects of Perceived Government Spending. Am. Econ. J.: Appl. Econ. 17 (4), 223–259.

Hallsworth, M., List, J.A., Metcalfe, R.D., Vlaev, I., 2017. The behavioralist as tax collector: Using natural field experiments to enhance tax compliance. J. Public Econ. 148.

Hanlon, M., Yeung, K., Zuo, L., 2022. Behavioral economics of accounting: A review of archival research on individual decision makers. Contemp. Account. Res. 39 (2), 1150–1214.

Hines, J.R., 2020. Income inequality, progressive taxation and tax expenditures.

Hoopes, J.L., Mescall, D., Pittman, J.A., 2012. Do IRS audits deter corporate tax avoidance? Account. Rev. 87 (5), 1603–1639.

Hoopes, J.L., Robinson, L., Slemrod, J., 2018. Public tax-return disclosure. J. Account. Econ. 66 (1), 142–162.

Huet-Vaughn, E., Robbett, A., Spitzer, M., 2019. A taste for taxes: Minimizing distortions using political preferences. J. Public Econ. 180, 104055.

IRS, 2020. Statistics of Income Tax Stats - County Data 2020.

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

Kahneman, D., Knetsch, J.L., Thaler, R.H., 1990. Experimental tests of the endowment effect and the coase theorem. J. Political Econ. 98 (6), 1325–1348.

Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. Econometrica 47 (2), 263–291.

Kerschbamer, R., Muller, D., 2020. Social preferences and political attitudes: An online experiment on a large heterogeneous sample. J. Public Econ. 182, 104076.

Lipkus, I.M., Samsa, G., Rimer, B.K., 2001. General performance on a numeracy scale among highly educated samples. Med. Decis. Making 21 (1), 37–44.

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

Luttmer, E.F.P., Singhal, M., 2014. Tax Morale. J. Econ. Perspect. 28 (4), 149–168.

Marshall, M.E., Farrar, J.M., Massey, D.W., Thorne, L., Wu, A., Bui, T., 2024. In all fairness: A meta-analysis of the tax fairness–tax compliance literature. Behav. Res. Account. 36 (2), 105–130.

Mayo, R., Moore, R., Ricks, K., 2021. Annual State and Local Government Finances Summary: 2019.

McKenzie, D., 2012. Beyond baseline and follow-up: The case for more T in experiments. J. Dev. Econ. 99 (2), 210–221.

Myles, G.D., Naylor, R.A., 1996. A model of tax evasion with group conformity and social customs. Eur. J. Political Econ. 12 (1), 49–66.

Nathan, B., Perez-Truglia, R., Zentner, A., 2025. My taxes are too Darn high: Why do households protest their taxes? Am. Econ. J.: Econ. Policy 17 (1), 273–310.

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

Peters, E., Dieckmann, N.F., Dixon, A., Hibbard, J.H., Mertz, C.K., 2006. Numeracy and decision making. Psychol. Sci. 17 (5), 407–413.

Sherlock, M.F., Marples, D.J., 2021. Overview of the Federal Tax System in 2022. Congr. Res. Serv..

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

Slemrod, J., 2019. Tax Compliance and Enforcement. J. Economic Lit. 57 (4).

Stantcheva, S., 2021. Understanding tax policy, how do people reason? Q. J. Econ. 136 (4), 2309–2369.

Tanner, R., Manwaring, P., 2023. Public Disclosure and Tax Compliance: Evidence from Uganda. The George Washington University, Institute for International Economic Policy.

Traxler, C., 2010. Social norms and conditional cooperative taxpayers. Eur. J. Political Econ. 26 (1), 89–103.

Tyran, J.-R., Sausgruber, R., 2006. A little fairness may induce a lot of redistribution in democracy. Eur. Econ. Rev. 50 (2), 469–485.

US Census Bureau, 2021. Population, Dallas County, Texas. https://www.Census.Gov/Quickfacts/Fact/Table/Dallascountytexas/POP010220.

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