Does Tax Avoidance Trickle Down?

Evidence from a Field Experiment

Justin Holz
University of Michigan

Ricardo Perez-Truglia
University of California, Los Angeles

Alejandro Zentner
University of Texas at Dallas

Abstract

The wealthiest individuals frequently engage in tax avoidance, and public awareness of this behavior is growing. Such awareness may undermine tax compliance among the broader population, either by revealing strategies others can emulate or by shifting social norms to make avoidance seem more acceptable. We first document that individuals are significantly less tolerant of both tax avoidance and tax evasion when carried out by the wealthiest households. We next present results from a field experiment on property tax appeals, which—like other forms of tax avoidance—is used disproportionately by the wealthiest households. Providing information about the appeal rates of the richest-1% of households significantly increased individuals' perceptions of unfairness, but did not affect their own appeal decisions—measured via administrative records—nor did it affect their expected tax savings. Information about the prevalence of appeals among comparable households did not affect appeal rates or expected tax savings. By contrast, information about the expected financial gains from appealing significantly increased appeals. Taken together, our results show that while tax avoidance by the wealthiest households is widely condemned, it does not spur imitation among the broader public.

JEL Classification: C93, H26, Z13.
Keywords: tax avoidance, inequality, learning, social norms.

This Draft: February 6, 2026. First Draft: September 1, 2025. Holz: holzj@umich.edu. Perez-Truglia: ricardo.truglia@anderson.ucla.edu. Zentner: azentner@utdallas.edu. We are thankful for excellent comments from Jacob Goldin. This project was reviewed and approved in advance by the Institutional Review Board at The University of Texas at Dallas. The field experiment was pre-registered in the AEA RCT Registry (AEARCTR-0009298). To prevent contamination of the subject pool (e.g., that subjects could read about the hypotheses being tested), we posted the RCT pre-registration immediately after the deadline to file an appeal had passed, but before the post-treatment data was available for analysis. After the study is accepted for publication, we will share all the code and data through a public repository. Alexia Witthaus Vine provided excellent research assistance.

1 Introduction

The old adage "every law has a loophole" rings especially true when it comes to tax law. No matter how meticulously tax codes are crafted, there are always opportunities for individuals and corporations to reduce their tax burden through legal means. Unlike tax evasion, which involves illegal actions and carries a heavy stigma, tax avoidance is often perceived more favorably, and sometimes even seen as a sign of financial savvy. Moreover, whereas evasion is hidden and difficult to detect, avoidance operates in the open, much like a magician performing a trick in full view of the audience (Slemrod and Yitzhaki, 2002; Feldman and Slemrod, 2007).

The wealthiest individuals and corporations are both the most capable of and the most likely to engage in tax avoidance. For example, some reports suggest that, thanks to sophisticated accounting strategies, billionaires face an average tax rate significantly lower than that of the typical American (e.g., ProPublica, 2023). The ultra-wealthy can afford teams of tax advisors and accountants devoted to crafting tailored avoidance strategies—efforts that are regularly covered in the media. One vivid illustration of the open nature of tax avoidance came during the 2016 presidential debate between Hillary Clinton and Donald Trump. When Clinton implied that Trump had paid no federal income tax in certain years, he quickly interjected, "That makes me smart" (Taylor, 2016). In subsequent interviews, Trump further defended his actions, stating, "as a businessman and real estate developer, I have legally used the tax laws to my benefit," and, "I fight like hell to pay as little as possible" (CNN, 2015).

Growing public awareness of tax avoidance by the wealthiest may undermine tax compliance among the broader population. The behavior of the rich may serve as a reminder that others can also exploit similar—or different—loopholes, or it may shift social norms to make avoidance seem more acceptable. We designed a natural field experiment to test this hypothesis.

In our field experiment, we study a particular form of tax avoidance: property tax appeals, also known as tax protests. Filing an appeal is a legal means by which households can reduce the amount they pay in property taxes. Property taxes are based on the market value of a home, which, unless the home has been sold recently, can be difficult to determine precisely. Households can leverage the subjective nature of home valuations when filing a tax appeal to argue that their homes are worth less than the assessed value. A majority of these appeals are successful, and can save households hundreds of dollars per year (Nathan et al., 2025).

Property tax appeals provide an ideal setting for our research for three reasons. First, unlike other forms of tax avoidance—such as offshore sheltering—that are accessible only to the very wealthy, property tax appeals are available to all homeowners, regardless of

income or wealth. Second, like other forms of tax avoidance, property tax appeals are far more common among the wealthiest households. For example, only 8% of households with homes valued at $200,000–$249,000 filed an appeal in 2021 in Dallas County, compared to 49% of the richest-1% of households (with homes valued above $1.9 million). Third, because property tax records are publicly available and easily accessible, we can provide participants with accurate information about appeals and also track their tax avoidance behavior.

We conducted our experiment in Dallas County, Texas, due to the logistical advantages of conducting a field experiment in a single location. With a population of approximately 2.6 million in 2020, as reported by the U.S. Census Bureau, Dallas County is Texas's second-largest county and surpasses the population of 15 U.S. states. The county showcases significant diversity across various dimensions, including ethnicity, and boasts a balanced mix of Democrat and Republican supporters. Moreover, property tax appeals operate in much the same way in this context as they do in many others (Dobay et al., 2019; World Bank, 2019).

Our research design isolates two distinct mechanisms through which learning about the prevalence of tax avoidance might influence an individual's decision to engage in similar behavior: observational learning (Cai et al., 2009) and social norms (Deutchman et al., 2025). The observational learning channel posits that when individuals discover others are avoiding taxes, they may update their beliefs about the potential savings they themselves could achieve, thereby increasing the perceived benefit of tax avoidance. The social norms channel suggests that learning others are engaging in tax avoidance may lessen the personal guilt associated with such behavior, making individuals more likely to follow suit. Our experiment was specifically designed to disentangle these two causal mechanisms. Moreover, our research design allows to distinguish between vertical comparisons—whether individuals are influenced by the behavior of the wealthiest households—and horizontal comparisons—whether they are influenced by behavior of households with homes of similar values.

Our experimental design is structured as follows. We mailed invitation letters to a sample of households to participate in an online survey a few weeks before the property tax appeal filing deadline. The field survey had an embedded information-provision experiment. We linked field survey responses with administrative records from the county assessor's office at the household level. This integration enabled us to provide customized information treatments to each household. Additionally, it allowed us to measure how the information randomized in the field survey influenced their subsequent decisions to file a tax appeal.

The field survey is structured around three key beliefs. The first concerns the share of the richest-1% of households—with homes worth over $1.9 million—that file tax appeals. For brevity, we refer to this as the *richest-1% appeal rate.* The second belief focuses on the share of households that file an appeal among those with home values within the same $50,000

bracket as the respondent. Hereinafter, we refer to these similar households as *comparable households* and to this belief as the *comparable appeal rate.* The third belief addresses the expected annual property tax savings from filing an appeal, which we refer to as the *expected savings*. This belief serves a dual purpose: first, it allows us to isolate the observational learning channel; second, it provides a monetary benchmark for interpreting the magnitude of the effects of the other two beliefs.

By isolating variation in posterior beliefs that arises solely from randomly assigned information, we can credibly estimate the causal effect of these beliefs on the decision to file a tax appeal. The experiment is structured as follows. The field survey begins by eliciting prior beliefs about the richest-1% appeal rate, the comparable appeal rate, and the expected savings. We then cross-randomize the provision of accurate feedback related to each belief. For instance, each respondent has a 50% chance of receiving information about the actual richest-1% appeal rate. Because the feedback is cross-randomized, respondents may receive information about all, some, or none of the three beliefs. After the information-provision stage, we elicit posterior beliefs.

We can measure how the information provision affects posterior beliefs, and also whether it affected the household's decision to file an appeal—measured using administrative records. Additionally, we draw on field outcomes to provide direct tests of the underlying causal mechanisms. The observational learning channel predicts that when households learn the richest-1% appeal more often, they may infer that appealing is a smart strategy and revise their expectations about their own potential tax savings. We can test this hypothesis by examining whether information about the richest-1% appeal rate affects respondents' expected savings. And the social norms channel predicts that learning the wealthiest appeal more often may increase the moral views on appealing. To test this hypothesis, we included a question eliciting the household's view on the acceptability of tax appeals.

To further examine individuals' attitudes toward tax appeals and other forms of tax avoidance, we conducted a complementary survey of 600 Americans, hereinafter referred to as the *Attitudes Survey.* We elicited perceptions of the acceptability of nine types of tax avoidance, including property tax appeals, as well as three forms of tax evasion. Respondents rated the general acceptability of these behaviors and their acceptability when undertaken by the richest-1% and the poorest-1%, separately. And to further examine individuals' attitudes, our field experiment included a revealed-preference measure of the acceptability of tax appeals. We inform participants that the researchers can assist a household with filing a tax appeal, and ask whether they wish to allow or prevent this assistance. We use a price-list method to elicit willingness to pay—either to enable or to block assistance—and measure it separately when the help is directed to a household in the richest-1% and when it is directed

to a household in the poorest-1%.

We mailed the invitations to the field experiment in April 2022, a few days after the window to file tax appeals had opened but around three weeks before it closed. We mailed invitations to 100,615 households, 3.9% of which responded to our field survey—a response rate that is slightly higher than benchmarks (Sinclair et al., 2012; Nathan et al., 2025). Our baseline sample consists of 3,681 respondents who completed the field survey and passed some basic filters. Before any adjustments from tax appeals, the average subject owned a home valued at $448,000 and had to pay $9,700 in property taxes for that year. Homeowners can choose to appeal their taxes directly—which is the primary outcome—but they can also hire an agent to appeal on their behalf too. For context, 45.5% of homeowners in the control group—those who answered the field survey but did not receive any informational interventions—appealed directly in 2022. The subject pool resembles the overall population of homeowners across many characteristics, but it overrepresents individuals with a higher propensity to file a tax appeal.

We begin by summarizing the results of the Attitudes Survey. We find that the three forms of tax evasion have the lowest acceptance ratings, averaging between 0.6 and 1.0 on a scale from 0 (completely unacceptable) to 3 (completely acceptable). In contrast, the acceptability of tax avoidance behaviors varies widely: profit shifting is viewed most similarly to evasion (with an average score of 1.3), while the most accepted form—avoiding consumption taxes through substitution—scores 2.4 on average. Property tax appeals ranks 4th among the nine types of tax avoidance, with a score of 2.3.

A second finding from the complementary Attitudes Survey is that each type of tax avoidance or evasion is consistently rated as substantially less acceptable when undertaken by someone in the richest-1% compared to someone in the poorest-1%. For instance, the average acceptability score for a property tax appeal is 1.9 when filed by a household in the richest-1%, compared to 2.5 when filed by a household in the poorest-1%. This wealth-based gap also emerges in our revealed-preference measure from the field experiment. A very small minority of subjects (4.2%) were willing to pay $30 to prevent helping a household in the poorest-1% file a tax appeal. In contrast, a much larger share (32.7%) were willing to pay $30 to prevent helping a household in the richest-1%.

Next, we present results from the information-provision experiment. We document substantial misperceptions across all three beliefs elicited in the field survey. Specifically, most households either underestimate or overestimate the average appeal rate of the richest-1%, the average appeal rate of comparable households, and the average savings of comparable households. This finding is consistent with a growing literature documenting limited information and bounded rationality among taxpayers (Feldman et al., 2016). We also show

that when provided with accurate information, households significantly update their posterior beliefs in the direction of the information—indicating that respondents were attentive and regarded the information as both relevant and credible. Lastly, we estimate the causal impact of beliefs on the likelihood of filing a tax appeal using a Two-Stages-Least-Squares (2SLS) model that leverages the exogenous variation in beliefs induced by the randomized information treatments (Cullen and Perez-Truglia, 2022).

We begin by summarizing the effects of the belief about expected savings, because it serves as natural benchmark for interpreting the rest of the results. We find that higher expected benefits from appealing have a positive, statistically significant, and economically meaningful effect on the subsequent probability of filing a tax appeal. Specifically, each additional $100 in expected savings raises the probability of appeal by 1.30 percentage points (pp), or 2.8% relative to the baseline rate. This result aligns with basic economic theory—individuals respond to financial incentives—and with prior quasi-experimental evidence (Nathan et al., 2025). This result is robust across model specifications and withstands multiple falsification checks—most notably, an event-study analysis around the date of the intervention.

We now turn to the effects of the belief concerning the richest-1% appeal rate. When individuals believe that the richest-1% appeal more often, perceived unfairness about the wealth gap in appeal rates rises significantly. This finding is consistent with the evidence noted above that people show less tolerance for avoidance among the wealthiest households. However, while it heightens the sense of unfairness, the belief about the appeal rate of the richest-1% does not significantly increase the household's own likelihood of filing a tax appeal. We find a coefficient that is negative, which is the opposite sign as predicted by the contagion hypothesis, close to zero (-0.034) and statistically insignificant. That is, a 1 pp increase in the perceived share of the richest-1% who file a tax appeal lowers the probability of filing a tax appeal by an insignificant 0.034 pp.

We argue that, based on the 90% confidence interval, this represents a relatively precise null. To support this, we benchmark the upper bound against the effects of the belief about expected savings and the treatment effects reported in other information-provision experiments conducted in the same context. However, we should not base conclusions on a single coefficient. Indeed, we test the mechanism in multiple ways such as leveraging alternative outcomes and the treatment arm for horizontal comparisons. All of this evidence, summarized below, reinforces the conclusion that perceptions about the richest-1% appeal rate have no effect on the decision to file a tax appeal.

For horizontal comparisons, the two underlying channels—observational learning and social norms—predict effects in the same direction as vertical comparisons. In other words, perceiving a higher appeal rate among comparable households should increase a household's

own likelihood of appealing. In fact, one may even expect horizontal comparisons to have stronger effects. Observational learning could be more influential when the observed behavior comes from households with comparable homes, to the extent that their appeals are perceived to better capture the individual's own costs and benefits from appealing. Moreover, to the extent that individuals compare themselves more with peers than with distant groups (Festinger, 1954; Clark and Senik, 2010), social norms may be shaped more strongly by the behavior of comparable households than the richest-1%. Yet, we find that the estimated effects of the perceived appeal rate of comparable homes are close to zero, statistically insignificant, and precisely estimated. The absence of contagion from comparable homes thus reinforces the conclusion that there is also no contagion from the richest-1%.

Using alternative outcomes, we can provide more direct tests for the observational learning and social norms channels. Against the prediction of the observational learning channel, we find that increasing perceptions about the richest-1% appeal rate does not affect a household's own expected tax savings—a null effect estimated with high precision. Likewise, learning that the comparable appeal rate is higher does not change an individual's own expected tax savings. For the social norms channel, we leverage a survey measure of injunctive norms—whether the respondent believes tax appeals are always justifiable. While the perceived appeal rate of the richest-1% does not influence a household's own moral view of tax appeals, we find mixed results for the perceived appeal rate of comparable households.

To evaluate how surprising our findings might be, we conducted a forecast survey with a sample of 84 experts, primarily professors with publications in related topics.[1] These experts received an explanation of the experiment and were asked to predict the causal effects of each of the three beliefs. A majority (66%) forecasted that beliefs about expected tax savings would positively influence the likelihood of filing an appeal, which aligns with our experimental results. There was an even stronger consensus among experts (70%) that the belief about the appeal rate of the richest-1% would have a positive effect on appeal rates, which contrasts with the null effects we find. And an even stronger majority (76%) predicted positive effects for the belief about the appeal rate of comparable homes, which also goes against our null finding. Furthermore, although experts predicted that beliefs about both the appeal rate of the richest-1% and that of comparable homes would have larger effects on the likelihood of filing a protest than beliefs about expected tax savings, we find the opposite.

We also discuss how our findings may generalize to other contexts. To the extent that property taxes and appeals operate similarly in other U.S. states and countries, the results may apply there as well. However, some limitations should be noted. Our sample overrepresents households predisposed to appeal, which could bias the results toward stronger effects.

[1] For more details about the design and results of the expert survey, see Appendix B.

For the null effects, though, this potential upward bias would only reinforce the conclusion. Generalizability may also vary by context. For example, we argue that our setting provides ideal conditions for the observational learning channel; therefore, the absence of evidence for this channel here suggests it is likely insignificant elsewhere. On the other hand, since tax appeals tend to be considered acceptable, the social norms channel could play a larger role in other, more complex or stigmatized forms of tax avoidance.

Our study contributes to several strands of literature. First, it relates to the growing body of research on tax avoidance and evasion among the wealthy. A substantial literature documents how the richest individuals employ a wide range of strategies to avoid taxes. For example, Smith et al. (2019) shows that the use of tax-preferred corporate forms is heavily concentrated at the top of the income distribution. Alstadsæter et al. (2019) shows that offshore tax evasion is overwhelmingly driven by the ultra-wealthy. Kleven et al. (2013) demonstrates that top earners engage in cross-border migration to minimize their tax burden. As a result of increased tax avoidance at the top of the distribution, effective tax rates—despite progressive tax schedules—end up being far less progressive than intended. In fact, some estimates suggest that the ultra-rich pay lower effective tax rates than the average American (ProPublica, 2023; Zucman, 2024).

While abundant evidence documents the extent of tax avoidance among the wealthiest individuals, our contribution is to examine whether this behavior affects tax compliance in the broader population. We show that the public is broadly aware of the higher prevalence of tax appeals among the richest-1% of households and generally disapproves of it. Yet, our results indicate that such perceptions do not translate into higher levels of tax avoidance in the general population.

Our study also relates to and contributes to the literature on network effects in tax compliance. A growing body of evidence shows that tax evasion choices can spread through word of mouth within social networks—primarily via tax preparers (Chetty et al., 2013; Boning et al., 2020; Battaglini et al., 2020; Brown et al., 2025), but also through other channels such as neighbors (Drago et al., 2020; Cruces et al., 2022), relatives (Alstadsæter et al., 2019), and coworkers (Bergolo et al., 2020).[2] The existing work focuses on word-of-mouth transmission. For instance, sending a letter with a tax enforcement threat to one taxpayer has been shown to increase compliance among their untreated neighbors (Drago et al., 2020; Cruces et al., 2022). The standard interpretation is that the recipient warns their neighbors about the audit threat, thereby influencing the neighbors' compliance behavior. Similarly, the literature argues that individuals may receive tax evasion or avoidance tips from their

[2] There is also evidence of related network effects that, while not directly involving tax evasion, pertain to taxation. For example, Wilson (2022) shows that earned income tax credit claiming behavior can spread through Facebook friend networks.

accountants, relatives, or coworkers. We contribute to this literature by examining two causal mechanisms—observational learning and social norms—that remain underexplored. Specifically, we hypothesize that, even absent word-of-mouth transmission, the mere awareness of others' tax avoidance can shape an individual's own decision to avoid taxes, either by signaling that avoidance is beneficial or by shifting the individual's moral views.[3]

The rest of the paper proceeds as follows. Section 2 describes the institutional context. Section 3 describes the design and implementation of the field experiment and the Attitudes Survey. Section 4 presents the results on attitudes towards tax avoidance. Sections 5 and 6 presents the results of the field experiment. The last section concludes.

2 Institutional Context

2.1 Dallas County

We conducted the field experiment in Dallas County, one of the largest counties in the United States, encompassing 17 cities with a combined population of about 2.6 million in 2022 (U.S. Census Bureau, 2021). Dallas is ethnically diverse, relatively balanced in political orientation (Nathan et al., 2025), and, like much of the country, characterized by pronounced inequality in home values—home to both several billionaires and a wide swath of low-income households. We use administrative data provided by the Dallas Central Appraisal District (DCAD).[4] The data from the DCAD includes information on ownership (e.g., name and address of the owner), property characteristics, and current and historic tax information including property taxable value, tax amounts and history of tax appeals.

2.2 Property Taxes

Property taxes fund various public services provided by its various jurisdictions (e.g., cities or school districts within the county) and the county itself (e.g., community colleges). The tax amount depends on the assessed property value, jurisdictional tax rates, and applicable household exemptions. Most importantly, households can apply to classify their primary residence as a homestead which, among other benefits, limits annual increases in assessed

[3] More broadly, our study relates to the growing literature on the drivers of tax morale; see Luttmer and Singhal (2014) and Slemrod (2019) for reviews. Recent examples include Cullen et al. (2021), who show that compliance is higher when taxpayers trust the government in power; Giaccobasso et al. (2025), who find that compliance rises when taxpayers believe they benefit from government spending; and Nathan et al. (2023), Ajzenman et al. (2024), and Best et al. (2025), who examine how fairness perceptions affect compliance.

[4] Appraisal districts in Texas are a political subdivision from the State of Texas responsible from appraising property for the purpose of taxable value assessment.

value to 10%. This is by far the most common exemption, with the vast majority of subjects in our sample claiming it.[5]

2.3 Tax Appeals

The DCAD assesses the value of all taxable property within the county as of January 1 each year, resulting in what is known as the *proposed value*. These proposed values are made publicly available in mid-April on the DCAD's website and are also mailed to some homeowners.[6] One important feature of this setting is the subjectivity involved in valuing a home. For recently sold properties, the sales price provides a clear benchmark. However, for homes that have not sold in some time, estimating their value requires some guesswork. Assessor offices often use statistical models—similar to those used by private companies like Zillow or Redfin—to determine assessed values. In fact, rather than conducting the analysis in-house, many assessor offices hire firms that specialize in real estate data. These models, however, have limitations—for example, they do not account for every observable characteristic of a home. A clear indication of their imprecision is the frequent discrepancy between different estimates. For example, Zillow and Redfin often provide significantly different valuations for the same home (Nathan et al., 2025); and homes often sell for more or less than what these models predict. As a result, proposed values are inherently ambiguous and subjective, making them vulnerable to challenges by homeowners seeking to reduce their tax burden.

Once homeowners are notified of the latest proposed value, they have one month to decide whether to file an appeal. Homeowners can choose to appeal on their own, either by filing an online form or by mailing a physical form.[7] At the end of our survey, we included an open-ended question about why people choose, or choose not to, appeal. Most subjects (53.3%) explicitly stated that their taxes were too high or that they wished to reduce their tax burden. This finding is consistent with other evidence (see Giaccobasso et al. (2025)) and supports anecdotal accounts that tax appeals are used as a form of tax avoidance, rather than solely to correct errors.

Homeowners can also choose to hire an agent to appeal on their behalf.[8] When appealing,

[5] There are other exemptions, that freeze the portion of property taxes allocated to school funding, such as the exemption for individuals older than 65 and the disability exemption, but they are uncommon in our sample (2.27% and 1.65% of subjects, respectively).

[6] Notifications are sent to households with increases in proposed values, changes in ownership, or other relevant updates. For a sample of the official notification letter, see Appendix H in Giaccobasso et al. (2025).

[7] Nathan et al. (2025) shows that in 2020, about 75% of direct appeals were filed online and 25% by mail.

[8] There is an industry of agents who help with property tax appeals in Texas. These agents charge flat fees or a percentage of the tax savings.

homeowners may argue that the proposed value exceeds the market value of their home.[9] Homeowners can share supporting evidence for their appeal, and they can also share their "Opinion of Value" which is how much they believe their property is truly worth. The DCAD then reviews the arguments for the appeal and may offer a settlement reducing the proposed value by email or phone. If the homeowner does not accept the settlement offer or the DCAD does not offer a settlement, the appeal moves to a hearing—for more details, see Nathan et al. (2025).

Since property taxes are calculated by applying the tax rate to a home's assessed value, one would expect that lowering the assessed value would immediately reduce the tax bill. However, for households with a homestead exemption, tax savings in the first year are not guaranteed. Under the Texas property code, the assessed value of a homestead property cannot increase by more than 10% per year—a threshold commonly referred to as the homestead cap. This cap can create a sharp discontinuity in the expected marginal tax savings from filing an appeal. For households without a homestead cap, a successful appeal that lowers the proposed value directly reduces the taxes owed. For households with homestead status, the effect depends on whether the cap is binding. If the homestead cap is not binding—the proposed value falls below the homestead cap value—any marginal reduction in the proposed value will result in a corresponding reduction in the tax bill. But when the homestead cap is binding, a marginal reduction in the proposed value will have no impact on the amount of taxes owed during the first year.[10]

As in other forms of tax avoidance (Landier and Plantin, 2017; Tyger and Eastman, 2019; Alstadsæter et al., 2022), tax appeals are significantly more prevalent among wealthier individuals. For example, the average appeal rate is 8% for a typical household (with a home valued $200,000–$249,000), compared to 49% for households in the richest-1% (with homes valued over $1.9 million). This disparity may stem from several factors. First, the expected financial benefit from tax avoidance is greater for high-income households, giving them stronger incentives to engage in such behavior.[11] Second, wealthier individuals tend to be more educated and financially sophisticated, which may make them more aware of tax appeals and better equipped to navigate them. Third, access to tax agents may play a significant role, as high-income households are more likely to use them.[12]

[9] Homeowners may also argue that comparable houses in their district have been assigned lower proposed values, resulting in unfair taxation. More rarely, they appeal arguing the existence of errors in property characteristics (e.g., incorrect square footage).

[10] For more details about how the homestead cap works, see Nathan et al. (2025).

[11] Conversely, the opportunity cost of time associated with managing tax appeals could be greater for wealthier households too.

[12] For instance, 2.12% of houses in the 50th percentile in 2022 appealed using an agent, while 10.7% of households in the richest-1% appealed using an agent.

3 Experimental Design and Implementation

3.1 Subject Recruitment

We mailed recruitment letters timed to arrive just as Dallas County homeowners became eligible to file tax appeals.[13] Each letter contained a unique survey code and a URL that recipients could use to participate in the study.[14] The unique survey code allowed us to personalize their field survey and link their responses to administrative property tax records. To incentivize participation, we promised subjects we would provide step-by-step instructions on filing appeals online upon completing the field survey.[15] In addition to serving as an incentive to participate, providing all subjects with instructions on how to appeal helps level the playing field: whether subjects appeal or not will depend more on whether they *want* to appeal and less on whether they know how to.[16]

We provide copies of the envelope in Appendix C and the letter in Appendix D. These materials included several features designed to increase the legitimacy of our study. First, we sent the letters on behalf of researchers at the University of Dallas, a well-known institution in Dallas County. Second, the envelope featured the logo of the University of Dallas, the name of a professor from that university, and non-profit organization postage. Third, the letter included a physical address for the researcher and a link to the study's website.[17] Fourth, we provided contact information for the researchers and the Institutional Review Board. Fifth, the letter and the envelope included each recipient's name and address. Finally, the letter also mentioned the recipient household's proposed value and estimated property tax amount for 2022.

3.2 Field Survey Design

In this section, we summarize the main features of the field survey. The full field survey instrument is attached as Appendix F. The main module can be split in three steps, which are summarized as follows:

[13] The letters could not be sent earlier, as we needed to wait for DCAD to release the final assessment values.

[14] When a property was jointly owned by multiple individuals, usually spouses, we mailed one letter listing all owners.

[15] This walkthrough was the same as the one provided in Appendix D of Nathan et al. (2025), with minor updates for the year 2022. Nathan et al. (2025) provides evidence that such instructions can significantly increase the likelihood that a household files a tax appeal.

[16] Some of our survey respondents may have participated in the survey conducted a year earlier by Giaccobasso et al. (2025). However, that earlier survey did not provide information about appeal rates among the richest-1%, comparable homes, or expected tax savings. Even if it had, any effects from the earlier study should be balanced across treatments.

[17] See Appendix E for a screenshot of the website.

- **Step 1:** We elicit the three prior beliefs. The first belief is about the richest-1% appeal rate in the previous year (2021). To make sure the subjects understand the question, we explain that the richest-1% are defined as homes valued at $1.9 million or more, and we provide pictures of a sample of three homes from this group (reproduced in Panel A of Figure 1). The second belief concerns the comparable appeal rate in the previous year (2021). We defined comparable homes as those assessed by the county within the same $50,000 value range. For example, homes comparable to one assessed at $233,000 would be those assessed between $200,000 and $249,999. The third belief pertains to the expected tax savings (in dollars) in the previous year (2021).

- **Step 2:** In the information-provision stage, we randomize information related to each of the three beliefs. More precisely, we use the administrative data to calculate the actual richest-1% appeal rate in 2021, the comparable appeal rate in 2021, and the expected savings in 2021. We cross-randomize subjects to receive each of the three pieces of information with a 50% probability for each. In other words, individuals are randomly assigned to one of eight treatment groups: one-eight receives no information (which we refer to as the pure control group), one-eight receive information about the expected savings only, ..., and one-eight receives all three pieces of information. We make the randomization explicit to subjects to avoid spurious updating. Specifically, we informed them that some participants would be randomly selected to receive information while others would not, and then either provided the information or notified them that they would not receive any.

- **Step 3:** We elicit the three posterior beliefs. We present all subjects with three questions similar to those in Step 1. Following Cavallo et al. (2017), and to avoid asking the exact same question twice, we frame the prior beliefs and feedback around the previous year (2021) and the posterior beliefs around the current year (2022). For the first posterior belief, we ask subjects to forecast the richest-1% appeal rate in 2022. For the second posterior belief, we ask them to forecast the comparable appeal rate in 2022. For the third posterior belief, we ask subjects to forecast the tax savings they would personally receive if they filed a tax appeal in 2022. We focus on their own expected savings because this is the most relevant factor for their decision. For example, a household may acknowledge that others achieve high tax savings from filing an appeal, yet still anticipate low savings for itself due to special circumstances—such as a binding homestead cap.

The goal of the experiment is to document whether the information shocks induced by the experiment affected posterior beliefs, and most importantly, whether the shocks subsequently

affected the probability of filing a tax appeal, measured via administrative records.

The effect of the information treatments may depend critically on the subject's prior beliefs. Subjects with accurate beliefs should not update their beliefs when they are provided with accurate information. Those who underestimate their expected savings or the appeal rate of other households may update their beliefs upward, while those who overestimate may adjust their beliefs downward. Thus, providing information to subjects will have heterogeneous effects depending on the individual's priors. By measuring prior beliefs before the information-provision stage, our research design allows us to accommodate heterogeneous updating. Following other studies (see, e.g., Cullen and Perez-Truglia, 2022; Giaccobasso et al., 2025), we employ a Two-Stages-Least-Squares (2SLS) model that estimates the causal effect of beliefs while allowing for heterogeneity by prior beliefs.

The main outcome of interest is whether the household subsequently filed an appeal, as recorded in administrative data. In addition, we included a few survey questions—elicited after the information-provision stage—that can serve as secondary outcomes aimed at probing the underlying causal mechanisms. Respondents were asked how likely they were to file an appeal this year, on a 1–4 likelihood scale. This outcome may be able to pick up short-term effects on the intention to appeal, even if those intentions do not translate into actual appeals. We also asked respondents to evaluate the fairness of the observed wealth gap in appeal rates. Specifically, the question restates the subject's posterior beliefs about the richest-1% appeal rate and the comparable appeal rate, and asks whether they view this as fair or unfair on a scale from 0 (very unfair) to 10 (very fair). The goal of this outcome is to assess whether the information treatments shifted perceptions of fairness. A final survey question asks, right after asking whether the household plans on filing a tax appeal, how justifiable it is to, "lawfully reduce your tax bill if you have a chance." This question was designed to probe the social norms mechanism, which posits that learning others are appealing may make such behavior feel more justifiable.[18] Lastly, the end of the field survey included a battery of questions on demographics such as gender, age, ethnicity, education level, and political party.

[18] There are a few caveats regarding the question used in the field experiment. First, we used the term "lawfully" when describing the behavior, which may have led respondents to view it as more acceptable. Second, we did not explicitly mention tax appeals, only stating "to lawfully reduce your tax bill if you have a chance." However, because this question immediately followed one asking whether the respondent planned to file a tax appeal in the coming year, we assume most subjects understood we were referring to tax appeals given the context.

3.3 Revealed-Preference Measure of Acceptability of Tax Appeals

One module of the field survey included a revealed-preference measure of the acceptability of tax appeals. We build on prior research showing that, due to significant hassle costs, providing households with step-by-step instructions substantially increases the likelihood of filing an appeal (Nathan et al., 2025). To assess whether participants view appeals as acceptable, we informed them that researchers could assist a household with filing a tax appeal and asked whether they wished to allow or prevent this assistance.

Specifically, we used an iterative multiple price list (iMPL), implemented separately for a household in the richest-1% and one in the poorest-1%, randomizing the order of these two questions to avoid potential order effects. To ensure households understand the composition of each of these two groups, each question includes three pictures of actual homes from that group—for reference, Panel A of Figure 1 shows the sample pictures for the richest-1% homes, while Panel B shows the corresponding pictures for the poorest-1%.

We explained to subjects that the researchers were considering sending detailed information to a household from the group, enabling them to appeal effortlessly if they wished. In practice, the help we provided was the same as in the field experiment by Nathan et al. (2025), which was shown to significantly increase the likelihood of an appeal. In one question, we asked subjects whether they wanted the researchers to help or not help a property owner randomly selected from the richest-1%. The other question was identical, except it referred to a property owner randomly selected from the poorest-1% instead of the richest-1%. In each question, subjects first had to choose whether they wanted the researchers to provide or withhold help. If they chose to help, we used the iMPL to elicit how much they were willing to pay to *provide* assistance. If they chose to withhold help, we used the iMPL to elicit how much they were willing to pay to *prevent* the household from receiving assistance.

The iMPL consists of a sequence of paired choices presented dynamically (Holz et al., 2024). Consider, for example, a subject who indicates a desire to help a household in the poorest-1% file a tax appeal. The first choice is between providing that help or receiving a \$15 bonus. The second choice depends on the first. If the subject chooses to help the homeowner, the second decision is between helping or receiving a \$24 bonus. If the subject in the first decision instead chooses the bonus, the second decision is between helping or receiving a \$6 bonus. We continue increasing or decreasing the bonus amount based on previous responses until the difference between the final two amounts is \$3 or the subject reaches the maximum bonus amount of \$30. For subjects who prefer to withhold help, the structure is slightly different. The initial choice is between providing help along with a \$15 bonus, or receiving no bonus (and thereby withholding help). Subsequent paired choices adjust the bonus amount up or down following the same logic as above. The complete

decision trees for the iMPL method are shown in Figure A.1. For each of the two scenarios—one involving the poorest-1% and one involving the richest-1%—the resulting willingness to pay (WTP) falls into one of 22 intervals: $(-\infty, -30)$, $(-30, -27)$, ..., $[-3, 0]$, $[0, 3]$, ..., $(27, 30]$, and $(30, \infty)$—positive values indicate a willingness to pay to provide help, while negative values indicate a willingness to pay to withhold it.

To ensure consequentiality, we selected one in every 100 subjects and implemented one of their decisions at random. At the start of this module, we reminded subjects that they might be selected and that one of their decisions could be implemented. Because of this potential for real-world consequences—unlike typical surveys that rely on stated preferences—our design elicits revealed preferences. This is particularly important in our context, as concerns about fairness are often inflated when choices carry no personal cost (Levitt and List, 2007), presumably because of social desirability bias (Bursztyn and Jensen, 2017).[19]

3.4 Subject Pool

We mailed field survey invitations to 100,615 unique households. This sample was drawn from the universe of 430,175 non-commercial, owner-occupied residential properties in Dallas County that paid property taxes in both 2020 and 2021. To construct the sample, we applied a set of basic filters—for example, restricting to single-family homes and excluding households in the richest-1% and the poorest-1%—and then randomly selected 100,615 households from the remaining list.[20]

We timed our letters to arrive before the appeal deadline so that subjects would have time to respond to the information we provided. We created the letters immediately following Dallas County's posting of the 2022 proposed values on April 18, 2022, and mailed them shortly afterward. Subjects began to respond to our field survey and visit the study website on April 22. A total of 4,174 participants started the field survey, and 4,005 finished the key module—that is, they responded to all three posterior beliefs. This response rate of 3.9% is slightly higher than those reported in previous studies in this context (Nathan et al., 2025; Giaccobasso et al., 2025). On average, respondents took 9.7 minutes to complete the field survey. Toward the end of the field survey, we included an attention check similar to those used in prior work (Giaccobasso et al., 2025), which 91.8% of respondents passed—a high rate in general, and especially given that the check appeared at the very end, when respondent fatigue was likely at its peak.

[19] For example, in the context of tax preferences, Ajzenman et al. (2024) show that wealthy households express support for a more progressive tax schedule, but when real stakes are involved, they fail to put their money where their mouth is.

[20] We excluded the richest-1% because a key goal is to understand how other households respond to their behavior. We exclude the poorest-1% because one of the survey questions involved them as well.

As explained in the RCT pre-registration, we dropped responses that could not be excluded ex-ante because of data availability. We dropped 122 subjects who, according to the DCAD's records, had already filed an appeal before starting our field survey, meaning that it was too late for any information included in the field survey to affect their decisions to appeal. Similarly, we exclude 93 subjects who responded to the field survey after the deadline to file an appeal, for whom it was too late for any information included in the field survey to affect their appeal decisions.

We also exclude observations with outlier beliefs. Extreme prior beliefs are probably not true misperceptions but the result of misunderstanding the question, making a typo, or not paying close attention to the task at hand. To reduce sensitivity to outliers, we follow the standard practice in information-provision experiments and drop respondents with extreme prior misperceptions (see e.g., Fuster et al., 2022; Bottan and Perez-Truglia, 2025).[21] In our baseline specification, we use a conservative definition of outliers that excludes 116 subjects based on their prior misperceptions.[22] After applying these filters, 3,681 respondents remain, forming our main subject pool. Because these exclusions are based on pre-treatment variables (e.g., prior beliefs), they should not compromise the validity of the experimental variation.

3.5 Descriptive Statistics and Randomization Balance

Column (1) of Table 1 provides some descriptive information about the subject pool. The average age is 47, around 42% of respondents are female, 55% are White, and 49% (33%) self-identify as Democrat (Republican). The average subject has a home assessed in about $448,000 and pays a property tax amount of about $9,700. Around 87% of the subject pool received a homestead exemption.[23] Moreover, given the rapid growth in home prices during this period, it is not uncommon for homestead caps to be binding. Specifically, 66.21% of the households had a binding homestead cap—that is, they were granted a homestead exemption and their appraised value had increased by more than 10% since the previous year.

Our main outcome of interest is whether the household filed a direct appeal in 2022.[24] Approximately 45% of control households filed a direct tax appeal that year. Households

[21] Moreover, extreme prior misperceptions may introduce attenuation bias in the 2SLS estimates. For example, consider an individual with an approximately accurate prior belief who appears to have an extreme misperception due to a typo (e.g., omitting a zero) when answering the question. We would not expect this individual's choices to be influenced by the provision of accurate information, as they already knew the information to begin with.

[22] For more details on the classification of outliers, see Appendix A.6. And Appendix A.8 presents the results without excluding outlier misperceptions.

[23] The high prevalence of the homestead exemption arises mechanically from the fact that we invited only households with owner-occupied homes.

[24] The appeal variable is based on data downloaded from the DCAD website on July 19, 2022.

could also hire a tax representative to file an appeal on their behalf: an additional 7.4% of control households did so. As outlined in the RCT pre-registration, the focus of this paper is on direct appeals, since our intervention is unlikely to affect agent-filed appeals. Among other reasons, the information was provided to owners rather than agents, so we would not expect effects on agent appeals—for further discussion, see Nathan et al. (2025) and Giaccobasso et al. (2025).

Columns (2) through (9) of Table 1 display the average baseline characteristics of each treatment.[25] For each of these pre-treatment characteristics, column (10) reports the p-value of a test of the null hypothesis that the averages are the same across all treatment groups. Table 1 shows that, consistent with successful random assignment, the observable characteristics are balanced across treatments.

In terms of observable characteristics (e.g., home value, number of bedrooms, or tax rate), the subject pool is similar to the universe of households in the county. Differences between field survey respondents and non-respondents are statistically significant but small (see Appendix A.1). However, one significant difference is that, relative to the universe of households, respondents to the field survey are substantially more likely to file an appeal. This pattern is consistent with Nathan et al. (2025) and Giaccobasso et al. (2025), who used similar recruitment strategies in this setting. This pattern is most likely due to the design of our letter, which describes tax appeals, so subjects considering filing an appeal in 2022 are likely to pay attention to the letter and thus also likely to notice the survey link included in the letter. Indeed, our invitation letter promised participants that the end of the field survey included instructions on how to file an appeal. Moreover, by reducing procedural barriers, the appeal instructions provided at the end of the study may have increased the likelihood that field survey respondents ultimately filed an appeal (Nathan et al., 2025).

3.6 The Attitudes Survey

We designed the Attitudes Survey to provide additional evidence on the acceptability of tax avoidance and evasion. A sample of the survey instrument appears in Appendix G. To keep the survey of a reasonable length, each subject saw a battery of questions for each of five forms of tax behavior.[26] One of those behaviors is always property tax appeals. Another three are chosen at random from a set of eight different tax avoidance strategies: tax deferral through retirement accounts, buy-borrow-die, harvesting tax losses, real estate depreciation,

[25] Some characteristics, such as the gender of the respondent, are asked after the information-provision stage. However, treatment assignment should not affect these responses.

[26] Because of the random assignment, the same survey instrument does not include every possible behavior a subject might have been shown. The full list of behaviors and their definitions appears in Table A.4.

carried interest, corporate profit shifting, tax avoidance through substitution and duty free stores. And another behavior is chosen at random from three tax evasion strategies: under reporting income, paying workers under the table, and falsifying deductions or expenses.

The five tax scenarios are presented in random order. For each scenario, we begin with a simple explanation of how the tax behavior works, typically illustrated with a concrete example that is accessible even to those with no prior expertise. Subjects are then asked to rate the acceptability of the behavior on a 0–3 scale, ranging from completely unacceptable (0) to completely acceptable (3). After evaluating its general acceptability, we ask about its acceptability when the behavior is undertaken by a household in the poorest-1%, and then again when it is undertaken by a household in the richest-1%. Recognizing that visuals can enhance communication compared to text alone (Perez-Truglia and Yusof, 2024), we present three sample photos each of homes in the poorest-1% and richest-1% groups.[27] The survey concludes with some questions about the subject's familiarity with the forms of tax behavior, whether they engaged themselves in each of these behaviors, and a standard battery of demographic questions.

4 The Social Acceptability of Tax Avoidance

4.1 Results from the Attitudes Survey

We recruited 600 respondents for the Attitudes Survey via Prolific, on March 13, 2025.[28] We limited eligibility to U.S. homeowners.[29] A great majority (83.33%) of subjects found the questions easy to understand, and 99.33% of the subjects passed our attention check.[30]Figure 2 presents the key results. Panel A presents the results for acceptability of each of these behaviors in general.[31] The most important takeaway is that while tax evasion behaviors are on average less accepted than tax avoidance, there is a wide range of acceptability across the different forms of tax avoidance. In other words, while technically all forms of tax avoidance are legal, some of them are more acceptable than others. We find that the three forms of tax

[27] The six photos are of homes in Dallas County and are the same images used in the field experiment. While cost of living varies across geographic areas and these homes may not be representative elsewhere, it would have been impractical to present a different set of images depending on the respondent's location.

[28] For a comparison of average subject characteristics between this survey and the field experiment, see Appendix A.4.

[29] To be able to compare responses between Texas and the rest of the country, we over-sampled responses from Texas: we collected 157 responses from Texas and 443 from other states.

[30] Çelebi et al. (2025) discuss the challenges posed by AI agents responding to surveys and show that, although these agents may threaten survey quality, they currently pose only a limited risk on Prolific. Additionally, at the end of the survey, we included an additional question designed to detect respondents who might be using Artificial Intelligence—see Appendix A.4.

[31] In both panels, the tax behaviors are ordered by their mean general acceptability from Panel A.

evasion have the lowest acceptance ratings, averaging between 0.6 and 1.0 on a scale from 0 (completely unacceptable) to 3 (completely acceptable). Some forms of tax avoidance, particularly those employed by big corporations or the ultra-wealthy, have low acceptability that are not far from the acceptability of tax evasion: corporate profit shifting, carried interest, and buy, borrow, die have an acceptability in the range of 1.3–1.4. Property tax appeals ranks fourth out of the nine types of tax avoidance, with a score of 2.3, ranking below tax avoidance through substitution, purchases from duty free stores, and tax deferral through retirement accounts.

While Panel A of Figure 2 shows the overall acceptability of each form of behavior, Panel B presents the acceptability separately for cases in which the behavior is undertaken by a household in the poorest-1% versus the richest-1%. For some tax behaviors, these results should be interpreted with caution, as the question referring to the poorest-1% may be a stretch—certain behaviors, such as buy-borrow-die or profit shifting, are realistically beyond the reach of the poorest households. With that caveat in mind, the pattern is consistent across all forms of tax behavior: acceptability is lower when the behavior is undertaken by a household in the richest-1% than when it is undertaken by a household in the poorest-1%. Take, for example, property tax appeals—one of the behaviors for which the comparison is most meaningful, since both poor and wealthy households can, and often do, appeal their property taxes. For appeals, the average acceptability scores are 1.86 for the richest-1% versus 2.51 for the poorest-1%, a difference that is large (0.65 points) and statistically significant (p-value < 0.001). Moreover, this wealth-based gap in acceptability is not only directionally consistent but also quantitatively similar across all other forms of tax avoidance and evasion. In sum, this evidence suggests that individuals are less willing to tolerate tax avoidance when it is carried out by the rich than by the poor.

4.2 Revealed-Preference Evidence on Acceptability of Tax Appeals

The results from the Attitudes Survey presented above are subject to the usual caveats associated with survey data, such as social desirability bias (Edwards, 1957). For instance, respondents may claim to view tax avoidance by the rich as unfair simply because they believe that is the socially appropriate view to express, even if they privately hold a different opinion. To address these concerns, we next present revealed-preference evidence on the matter, using a novel approach from the field experiment. The magnitudes are not directly comparable between the Attitudes Survey and the field experiment due to differences in subject recruitment and methodology. Nevertheless, the results are qualitatively consistent:

there is strong support for appeals by households in the poorest-1%, but markedly less tolerance for appeals by households in the richest-1%.

The detailed results for the willingness to pay results are presented in Figure 3. Panel A displays households' willingness to provide or withhold help from another household in the poorest-1%, while Panel B displays the corresponding willingness to help a household in the richest-1%. In each panel, positive values indicate a willingness to pay to help someone else file a tax appeal, while negative values indicate a willingness to pay to prevent the help from occurring. The rightmost and leftmost bars are $\geq$ \$30 and $\leq -$\$30 by construction, because \$30 was the maximum reward amount offered to subjects in the experimental design. The middle category, \$0, corresponds to subjects who were approximately indifferent as to whether to provide or prevent the help—more precisely, cases in which the willingness to pay is less than \$3 in absolute value.

While there is high acceptance of tax appeals when filed by a household in the poorest-1% (Panel A of Figure 3), that acceptance is markedly lower when filed by a household from the richest-1% (Panel B). Specifically, while 82.4% of respondents were willing to pay \$30 to help a household from the poorest-1% file an appeal, that share drops sharply when it comes to helping the wealthy: only 30.5% were willing to pay the same amount to help a household from the richest-1% file an appeal ($p < 0.001$).[32] At the other end of the spectrum, whereas 4.0% of respondents were willing to pay \$30 or more to prevent help to a household in the *poorest-1%*, a substantially larger share (31.4%) were willing to pay that amount to prevent help to a household in the *richest-1%* ($p < 0.001$).

Lastly, Table A.5 shows broad acceptance across subgroups for supporting the poorest-1% in reducing their property tax burdens. The share favoring aid to the poorest-1% only ranges from 88.26% to 91.96%, while the share opposing such aid ranges from only 4.46% to 6.43%. These patterns indicate near-universal support for assistance to the lowest-income households, with little meaningful variation across observable characteristics.

In contrast, there is substantial heterogeneity in willingness to help the richest-1% file a tax appeal. The largest differences emerge between respondents who view appeal inequality as fair and those who view it as unfair. Those who find appeal inequality unfair are 20 pp (43%) less likely to favor helping the richest-1% appeal ($p < 0.001$) and 20 pp (57%) more likely to pay to prevent us from helping the richest-1% appeal ($p < 0.001$). Partisanship also emerges as a significant factor in support for assisting the richest-1%. While Republican and Democratic subjects are similarly likely to want to help households in the poorest-1% (91.96% vs. 90.06%), Republicans are significantly more likely to help households in the

[32] Our expert forecast survey also included two questions related to these questions on willingness to pay—see Appendix B for more details.

richest-1% (49.87% vs. 33.10%, $p < 0.001$). Finally, willingness to help either the richest-1% or the poorest-1% is largely unrelated to the subject's own wealth—the differences are small in magnitude and statistically insignificant.

5 Prior Beliefs and Belief Updating

5.1 Accuracy of Prior Beliefs

We find substantial misperceptions across all three beliefs elicited in the survey. We start with Panel A of Figure 4, which shows the distribution of misperceptions about the richest-1% appeal rates—that is, the difference between the share of richest-1% of households who actually filed a tax appeal in the previous year (i.e., the feedback provided in the experiment) and the subject's own prior belief about this share. For reference, the correct answer was 49%. We find large misperceptions. Only a small group (9.07%) guessed the appeal rate of the richest-1% within 5% of the truth. The mean absolute error shows that, on average, prior beliefs were off by 29.16 pp. Moreover, these misperceptions were somewhat skewed: on average, subjects overestimated the richest-1% appeal rate by 8.49 pp.

Panel A of Figure 5 shows the distribution of misperceptions about the comparable appeal rate. The correct answer to this question depends on the subject's home value bin—for instance, for homes valued at \$200,000–\$249,999, the correct answer is 8%.[33] There are large misperceptions about the comparable appeal rate too. By one metric—the wisdom of the crowds—individuals are less informed about the comparable appeal rate than about the richest-1% appeal rate. On average, subjects overestimate the comparable appeal rate by 19.65 pp, whereas they overestimate the richest-1% appeal rate by only 8.49 pp.

Furthermore, we can assess more directly whether people are aware of the gap in appeal rates between the richest-1% and comparable homes. The richest-1% are far more likely to file an appeal. For example, for the average respondent, the appeal rate is 49% for the richest-1% compared to 8% for comparable homes—a difference of 41 pp. On the one hand, households overestimate the richest-1% appeal rate by 7 pp, suggesting a general awareness that the wealthiest appeal at high rates. On the other hand, households also overestimate the comparable appeal rate by 18 pp. As a result, the average household underestimates the wealth gap in appeal rates by approximately 11 pp. In sum, although some households underestimate and others overestimate, on average they recognize that the richest-1% file appeals at much higher rates than comparable households.

One heuristic that typically does a good job at explaining misperceptions is extrapolation.

[33] Feedback for subjects in all other bins is provided in Table A.2.

Just to mention one example, when asked about the salary of co-workers, individuals tend to report their own salary as their best guess (Cullen and Perez-Truglia, 2022).[34] According to this heuristic, if subjects were extrapolating from their own situation, they should be better at guessing about comparable homes than about the richest-1%. The fact that guesses about the richest-1% are similarly accurate—or even more accurate—than guesses about comparable homes suggests that households have some access to information about the appeals of the richest-1%, such as through news coverage. Another possibility is that individuals inferred that the wealthiest appeal more often for various reasons, such as having more to gain from an appeal or greater access to professional assistance.

Panel A of Figure 6 shows the distribution of misperceptions about expected savings—that is, the difference between the *actual* savings for comparable households in the previous year (i.e., the feedback provided in the experiment) and the subject's prior belief. The feedback depends on the subject's home value bin—for instance, for homes valued at \$200,000–\$249,999, the feedback was \$322. For reference, these \$322 in savings represent about 6.34% of the average tax bill in this bracket.[35] The mean absolute error shows that, on average, prior beliefs were off by \$721.[36] Misperceptions were also skewed: on average, subjects overestimated the savings by \$591. That is, on average, subjects believed the potential savings were twice as large as they actually are.[37]

Since the question on expected savings is cognitively more demanding, we acknowledge that some of these misperceptions may arise from the complexity of the belief being elicited (Feldman et al., 2016). We were careful in how we worded the question and conducted small pilots to fine-tune the language. Despite these efforts, it is possible that some subjects misinterpreted the question. For example, some subjects might have responded about the expected savings conditional on a successful appeal, rather than the unconditional expectation we intended. Even so, the fact that a strong majority of subjects reported beliefs of the correct order of magnitude suggests that they understood the question.[38]

[34] There is evidence of this heuristic in a broad set of contexts such as perceived tax rates (Nathan et al., 2023) and perceived relative income (Cruces et al., 2013).

[35] Feedback for subjects in all other bins is provided in Table A.2.

[36] The results already exclude outliers. Including them would mechanically inflate the degree of misperception reported here. However, we believe that would be misleading, as those subjects likely did not understand the question in the first place.

[37] We caution that this overestimation may be less pronounced in the general population. If individuals who expect to save more from filing an appeal are more likely to respond to our survey, those who overestimate their savings would be mechanically overrepresented in our sample.

[38] A more significant misinterpretation would be to report the expected reduction in the proposed property value rather than the expected reduction in tax savings. In that case, their responses would be in a different order of magnitude—tens of thousands of dollars instead of hundreds. However, because such individuals would not even get the order of magnitude right, they would get excluded from the analysis under our definition of outliers—see Appendix A.5 for details.

5.2 Belief Updating

For each of the three beliefs we study, individuals update their beliefs substantially when provided with accurate information through the experiment. To quantify the degree of belief updating, we adopt the framework commonly used in other information-provision studies (e.g., Cavallo et al., 2017; Cullen and Perez-Truglia, 2022; Giaccobasso et al., 2025). Let subscript i denote individuals and subscript t denote the current year (2022) and $t-1$ denote the previous year (2021). Define $w_{i,t-1}^{prior}$ as individual i's prior belief about the richest-1% appeal rate in 2021. Let $w_{i,t-1}^{feed}$ be the piece of information being randomly assigned—the actual share of richest-1% households who filed an appeal in 2021. In our experiment, this variable does not require a subscript i because the feedback was not personalized—however, we retain the subscript i to accommodate the more general case in which it could be. Let $w_{i,t-1}^{post}$ denote the posterior belief about the richest-1% appeal rate in 2021, which we do not measure in the survey (to avoid asking the same question twice) but define here for the sake of the framework. And let $w_{i,t}^{post}$ be the posterior belief about the richest-1% appeal rate in 2022. Analogously, we define $c_{i,t-1}^{prior}, c_{i,t-1}^{feed}, c_{i,t-1}^{post}, c_{i,t}^{post}$ for beliefs about the comparable appeal rate and $s_{i,t-1}^{prior}, s_{i,t-1}^{feed}, s_{i,t-1}^{post}, s_{i,t}^{post}$ for beliefs about the expected savings.

We assume that firms form their future expectations by projecting their perceptions about the past. For example, in the case of the belief about the richest-1% appeal rate:

$$w_{i,t}^{post} = \mu + \theta^{w} \cdot w_{i,t-1}^{post}, \tag{1}$$

where $\theta^{w} > 0$ captures the degree of pass-through from past to future beliefs: that is, for each 1 pp increase in an individual's perceived 2021 appeal rate of the richest-1%, his or her expected 2022 appeal rate for the richest-1% rises by θ^{w} pp. For the beliefs about comparable appeal rate, θ^{c} would capture the extent to which individuals extrapolate the share of comparable homes that appealed in the past year to the share they expect will appeal in the following year. For expected savings, θ^{s} would measure the extent to which subjects extrapolate from the average savings of comparable homes in 2021 to their *own* expected savings in 2022—that is, extrapolating jointly from past to future and from others to themselves.

Now we turn to the Bayesian updating. Let T_i^{w} be an indicator variable equal to 1 if the individual was randomly chosen to receive the information about the richest-1% appeal rate in 2021 and 0 otherwise—for the other two beliefs, the corresponding indicator variables would be T_i^{c} and T_i^{s}. Think of the case when the individual receives the information ($T_i^{w} = 1$).

Based on the assumptions of a Bayesian learning model with Gaussian distributions,[39] after observing the information the individual is expected to update the posterior belief as follows:

$$w_{i,t-1}^{post} = \alpha^w \cdot w_{i,t-1}^{feed} + (1 - \alpha^w) \cdot w_{i,t-1}^{prior}, \tag{2}$$

where $\alpha^w \in [0, 1]$ is the weight assigned to the new information relative to the prior belief, which depends on the relative accuracy of the prior belief relative to the accuracy of the signal. Likewise, α^c and α^s would be the corresponding learning rates for the beliefs about comparable appeal rates and expected savings, respectively. If we combine equations (1) and (2) we obtain the following expression:[40]

$$\underbrace{w_{i,t}^{post} - w_{i,t-1}^{prior}}_{\text{Belief Update}} = \mu + \theta^w \cdot \alpha^w \cdot \underbrace{\left(w_{i,t-1}^{feed} - w_{i,t-1}^{prior}\right)}_{\text{Prior Gap}} + (\theta^w - 1) \cdot w_{i,t-1}^{prior} \tag{3}$$

The key prediction from the model is that, for individuals who received information, the belief updates should be a linear function of the prior gaps. Intuitively, respondents who overestimated the true richest-1% appeal rate should revise their beliefs downward upon receiving feedback; those who underestimated it should revise upward; and those who were already accurate should exhibit no updating. Note also that the strength of this belief updating is given by the product of the two key parameters: the learning weight (α^w) and the degree of extrapolation (θ^w).

There is one potential issue with estimating equation (3) directly. In practice, individuals may update their beliefs in the direction of the feedback for spurious reasons, that is, even if they did not get to see the information. Simply being asked the same question twice may lead them to think more carefully, reconsider their earlier response, or correct typographical errors, all of which could move their second answer closer to the truth. To disentangle true learning from spurious learning, we estimate the following specification that leverages the randomized nature of the information provision:

$$w_{i,t}^{post} - w_{i,t-1}^{prior} = \gamma_0 + \gamma_1 \cdot T_i^w \cdot \left(w_{i,t-1}^{feed} - w_{i,t-1}^{prior}\right) + \gamma_2 \cdot \left(w_{i,t-1}^{feed} - w_{i,t-1}^{prior}\right) + \gamma_3 \cdot w_{i,t-1}^{prior} + \epsilon_{i,t} \tag{4}$$

Intuitively, the key parameter is γ_1, which measures whether the slope between belief updates and prior gaps is stronger for individuals who received information ($T_i^w = 1$) than those who did not ($T_i^w = 0$).[41]

[39] These assumptions include that the priors and the signal are normally distributed, and the variance of the prior and the signal is independent of the prior's mean. See Section C of Cavallo et al. (2017) for further discussion.

[40] Specifically, equation (3) is obtained by plugging equation (2) into equation (1) and subtracting $w_{i,t-1}^{prior}$ from both sides.

[41] In practice, all subjects received identical feedback about the richest-1% appeal rate. Consequently, $(w_{i,t-1}^{feed} - w_{i,t-1}^{prior})$ and $w_{i,t-1}^{prior}$ are perfectly collinear, so controlling for one is sufficient.

Panel B of Figure 4 shows the updating results corresponding to the belief about the richest-1% appeal rate. For simplicity, this figure includes only individuals who either received no information or received information solely about the richest-1% appeal rate, thereby isolating the effect of that specific information. The x-axis shows the prior gaps, and the y-axis shows the belief updates. Even in the control group (who received no information), there is a positive and statistically significant association between belief updating and prior gaps, but small in magnitude. Indeed, pattern is commonly observed in information-provision experiments and typically attributed to spurious learning (e.g., Cavallo et al., 2017; Cullen and Perez-Truglia, 2022; Giaccobasso et al., 2025).[42]

The key takeaway from Panel B of Figure 4 is that genuine learning occurs: the association between belief revisions and prior gaps is much stronger among those who received the information (0.886) than among those who did not (0.177). The difference between the two slopes (0.709) is large and statistically significant. For simplicity, let's call the product of the treatment indicator and the prior gap the *information shock*. In a world where individuals fully trust the information ($\alpha^w = 1$) and fully extrapolate ($\theta^w = 1$), a 1 pp increase in the information shock would be expected to translate into a 1 pp belief update. In comparison, the difference between slopes implies that a 1 pp information shock leads to a 0.709 pp belief update, indicating that individuals react less than fully to the information shocks, though fairly close to it.[43]

We also find strong updating for the other two beliefs, though with somewhat different magnitudes. Panel B of Figure 5 shows the learning results for beliefs about the comparable appeal rate. The difference in slopes between the treatment and control groups is large (0.709) and statistically significant (p-value<0.001). This degree of belief updating is exactly the same magnitude as the corresponding degree of updating for the richest-1%. A natural interpretation is that individuals hold equally strong priors about comparable homes as they do about the richest 1%.

Panel B of Figure 6 presents the learning results for beliefs about expected savings. Consistent with genuine learning, the difference in slopes between the treatment and control groups is large (0.446) and statistically significant (p-value=0.002). The degree of belief updating for expected savings (0.446) is in the same order of magnitude but somewhat smaller than for the other two beliefs (0.709). This lower degree of updating has an intuitive explanation: when updating their beliefs, households are extrapolating not only from the

[42] However, in our application, because the prior and posterior beliefs refer to different periods, this pattern could be attributed to the extrapolation—specifically, a $\theta^w = 0.826$ would predict the observed slope of 0.174.

[43] While this is true on average, updating can be highly heterogeneous, with some individuals update fully but others not updating at all, presumably because they are not paying attention—see Cullen and Perez-Truglia (2022) for a detailed discussion.

past year to the current year, but also extrapolating about comparable homes to their own household. As a result, households facing a binding homestead cap should be expected to extrapolate much less—or not at all. To illustrate this point, consider two identical households, A and B, both of which underestimated the average savings of comparable homes by $100. After receiving accurate information, both households revise their posterior belief about the average tax savings of comparable homes upward by $100. The difference is that household A faces a binding homestead cap, meaning a marginal reduction in the proposed value has no effect on the tax amount they owe this year, whereas household B does not. As a result, even though both households learn that others save $100 more than they initially believed, household A may not update its own expected tax savings at all, or at least will update less strongly than household B. Consistent with this prediction, Appendix A.6 shows that the degree of belief updating is substantially stronger for households without a binding homestead cap than for those with one. Indeed, the degree of belief updating about tax savings among households without a binding homestead cap (0.696) is similar to the degree of updating documented above for the other two beliefs (0.709).

The evidence discussed above shows that when individuals receive information pertaining to one belief, they update the corresponding posterior belief. However, there could also be *cross-learning*, whereby information related to one belief influences other beliefs as well. We find limited evidence of cross-learning, suggesting that belief updating was largely compartmentalized—the corresponding results are presented in Section 6.7 and Appendix A.6.[44]

6 Causal Effects of Information

6.1 Total Effects vs. Partial Effects

We aim to estimate the causal effects of each of the three posterior beliefs on the decision to file a tax appeal and other outcomes. To do so, we follow the same 2SLS approach used in other information-provision experiments (e.g., Cullen and Perez-Truglia, 2022; Giaccobasso et al., 2025), exploiting the exogenous variation in posterior beliefs generated by the randomized provision of information. There is, however, one methodological distinction from prior studies. In our research design, we expect the three beliefs to be potentially interrelated. For instance, when individuals receive information about the richest-1% appeal rate, they may update their posterior belief about that rate but also consider the information sufficiently

[44] In two cases we find effects that are statistically significant, but quantitatively small: when individuals update their posterior beliefs about the richest-1% appeal rate, they slightly adjust their beliefs regarding the comparable appeal rates, and vice versa—see Appendix A.6 for more details.

relevant to revise their beliefs about their own expected tax savings. In other words, while other studies often treat cross-learning as a nuisance, in our design it plays a central role in understanding the causal mechanisms.

We distinguish between two effects of information, which we term the "total effect" and the "partial effect," drawing an analogy to total and partial derivatives. The total effect of a piece of information encompasses both the impact through the direct belief plus any indirect effects operating through other beliefs. The partial effect, by contrast, isolates the direct impact after controlling for the indirect effects through other beliefs. For instance, suppose that after receiving accurate information about the richest-1% appeal rate, individuals change their choices partly because the information directly affects their belief about the richest-1% and partly because it indirectly alters their belief about their own expected savings. The partial effect captures only the direct impact on the belief about the richest-1%, while the total effect captures both the direct and indirect effects.

6.2 Econometric Model: Total Effects

We begin our analysis with the total effect of the belief concerning the richest-1% appeal rate. Our main outcome of interest, $P_{i,t}$, is an indicator variable equal to 100 for individuals who filed an appeal in 2022 (i.e., post-treatment) and 0 otherwise. While the following model is applicable to each of the three beliefs, for the sake of brevity, we focus on the belief about the richest-1% appeal rate. We seek to estimate a regression of $P_{i,t}$ on the posterior belief $w_{i,t}^{post}$. However, because posterior beliefs may be endogenous, such a regression would not necessarily identify the causal effect of that belief. To achieve causal identification, we employ a 2SLS model that exploits only the exogenous variation in posterior beliefs induced by the randomization of information. For the total effect, the relevant subsample consists of individuals in two treatment groups: those who received information solely about the appeal rate of the richest-1% and those who did not receive any information. The specification of interest is:

$$P_{i,t} = \beta_0 + \beta_w^{total} \cdot w_{i,t}^{post} + \beta_2 \cdot (w_{i,t-1}^{feed} - w_{i,t-1}^{prior}) + \beta_3 \cdot w_{i,t-1}^{prior} + X_i\beta_X + \nu_i \tag{5}$$

The endogenous variable is $w_{i,t}^{post}$, and the excluded instrument is $T_i^w \cdot (w_{i,t-1}^{feed} - w_{i,t-1}^{prior})$.[45] The vector X_i corresponds to a set of additional controls, which in the baseline specification

[45] Formally, the exogeneity assumption is $\mathbb{E}\left[\left(w_{i,t-1}^{feed} - w_{i,t-1}^{prior}\right) \cdot T_i^w \cdot \nu_i \mid W_i\right] = 0$ where W_i is the vector of control variables $\{w_{i,t-1}^{feed} - w_{i,t-1}^{prior}, w_{i,t-1}^{prior}, X_i\}$. In plain terms, to interpret the 2SLS estimate as the causal effect of the belief, we must assume that heterogeneity in the effects of information is driven solely by differences in prior misperceptions, and not by unobserved factors correlated with those misperceptions.

corresponds to the priors and gaps of the other two beliefs.[46]

We can illustrate the intuition behind the 2SLS model with a simple example. This model is based on the comparison between pairs of individuals with the same prior gap. For example, consider two individuals who both underestimate the richest-1% appeal rate by 10 pp. We then randomly assign information about the true richest-1% appeal rate to one of them. We would expect that, relative to the individual who does not receive the information, the individual who does will end up with a higher posterior belief about the richest-1% appeal rate. For the sake of argument, suppose the individual who did not receive the information continues to underestimate the actual richest-1% appeal rate by 10 pp, while the informed individual reacts strongly and now underestimates by only 2 pp. The information provision can thus be interpreted as a positive shock of 8 pp to the perceived richest-1% appeal rate. We can then examine how this 8 pp shift in posterior beliefs translates into differences in behavior—specifically, whether it leads to a higher probability of filing an appeal.[47]

We can estimate the same model described above for the other two beliefs, using a specification analogous to equation (5) but with the corresponding variables for each belief.[48] When estimating the model for the belief about expected savings, there is one small difference in the specification. As discussed in Section 2 above, whether a household faces a binding homestead cap can play a major role in shaping how it updates its expectations about its own tax savings. Thus, allowing for this heterogeneity in the first stage can improve the asymptotic mean squared error of the 2SLS estimator (Abadie et al., 2024). Specifically, let H_i be an indicator variable for whether household i has a binding homestead cap. We follow the fully-interacted method from Abadie et al. (2024), allowing for interactions of H_i with the excluded instrument.[49]

Our first hypothesis is that $\beta_w^{total} > 0$: if individuals increase the perceived richest-1% appeal rate, they will be more likely to file an appeal themselves, potentially due to social norms or observational learning. Similarly, we hypothesize that $\beta_c^{total} > 0$: if individuals

[46] In practice, the feedback about the appeal rate of the richest-1% was identical for all subjects. As a result, $(w_{i,t-1}^{feed} - w_{i,t-1}^{prior})$ and $w_{i,t-1}^{prior}$ are perfectly collinear, so it suffices to control for only one of them.

[47] For further discussion on this 2SLS model, see (Cullen and Perez-Truglia, 2022). One caveat of the 2SLS estimate is that it identifies a Local Average Treatment Effect (LATE). In other words, it captures the average effect of beliefs among individuals whose posterior beliefs changed as a result of the information experiment. Thus, by construction, the 2SLS estimate places greater weight on individuals with larger prior misperceptions and, conditional on those misperceptions, on those who respond more strongly to the feedback.

[48] Each 2SLS model is estimated with a different pair of treatment groups. For instance, the model for the belief about the comparable appeal rate is estimated using individuals who either did not receive any information or received only information about the comparable appeal rate.

[49] More precisely, the endogenous variable is $s_{i,t}^{post}$, the excluded instruments are $H_i \cdot T_i^s \cdot (s_{i,t-1}^{feed} - s_{i,t-1}^{prior})$, $(1 - H_i) \cdot T_i^s \cdot (s_{i,t-1}^{feed} - s_{i,t-1}^{prior})$ and we control for $H_i \cdot (s_{i,t-1}^{feed} - s_{i,t-1}^{prior})$, $(1 - H_i) \cdot (s_{i,t-1}^{feed} - s_{i,t-1}^{prior})$, $H_i \cdot s_{i,t-1}^{prior}$, $(1 - H_i) \cdot s_{i,t-1}^{prior}$ and H_i.

increase the perceived comparable appeal rate, they will also be more likely to file an appeal, again possibly due to social norms or observational learning. Finally, we hypothesize that $\beta_s^{total} > 0$: increasing a household's expected savings will make them more likely to appeal, as individuals respond to pecuniary incentives. That is, if they expect an action to yield larger financial gains, they are more likely to take it.

6.3 Econometric Model: Partial Effects

To disentangle the causal mechanisms at play, we next focus on estimating the partial effects. For this purpose, we estimate a single regression that pools all eight treatment groups and includes all three beliefs simultaneously as right-hand-side variables:

$$\begin{aligned} P_{i,t} =& \beta_0 + \beta_w^{partial} \cdot w_{i,t}^{post} + \beta_c^{partial} \cdot c_{i,t}^{post} + \beta_s^{partial} \cdot s_{i,t}^{post} + \\ &+ \beta_1 \cdot (w_{i,t-1}^{feed} - w_{i,t-1}^{prior}) + \beta_2 \cdot (c_{i,t-1}^{feed} - c_{i,t-1}^{prior}) \\ &+ \beta_3 \cdot H_i \cdot (s_{i,t-1}^{feed} - s_{i,t-1}^{prior}) + \beta_4 \cdot (1 - H_i) \cdot (s_{i,t-1}^{feed} - s_{i,t-1}^{prior}) + \\ &+ \beta_5 \cdot w_{i,t-1}^{prior} + \beta_6 \cdot c_{i,t-1}^{prior} + \beta_7 \cdot H_i \cdot s_{i,t-1}^{prior} + \beta_8 \cdot (1 - H_i) \cdot s_{i,t-1}^{prior} + \\ &+ \beta_9 \cdot H_i + X_i \beta_X + \nu_i \end{aligned} \tag{6}$$

The endogenous variables are $w_{i,t}^{post}$, $c_{i,t}^{post}$ and $s_{i,t}^{post}$, and the excluded instruments are $T_i^w \cdot (w_{i,t-1}^{feed} - w_{i,t-1}^{prior})$, $T_i^c \cdot (c_{i,t-1}^{feed} - c_{i,t-1}^{prior})$, $H_i \cdot T_i^s \cdot (s_{i,t-1}^{feed} - s_{i,t-1}^{prior})$, and $(1 - H_i) \cdot T_i^s \cdot (s_{i,t-1}^{feed} - s_{i,t-1}^{prior})$. The coefficients $\beta_w^{partial}$, $\beta_c^{partial}$ and $\beta_s^{partial}$ are the three partial effects of interest. For example, $\beta_w^{partial}$ measures the effect of the belief about the richest-1% appeal rate while holding constant the other two beliefs—about the comparable appeal rate and expected savings. Relative to the total effects, the partial effects have the advantage of pinpointing the precise causal mechanisms at play.[50]

Ex ante, whether total and partial effects differ depends on how individuals process information. The two effects will coincide if learning is *compartmentalized*—that is, when individuals update only the belief directly targeted by the information and ignore it when forming other beliefs. Ex post, as shown later, learning appears highly compartmentalized, so we would expect the partial and total effects to be similar. However, another possible source of divergence between total and partial effects is *information saturation*. In estimating the total effect, we compare individuals who received a single piece of information to those who received none. By contrast, the partial effect is identified among individuals who may have received two or three pieces of information simultaneously. If attention is limited, the impact

[50] As noted earlier, because the feedback about the appeal rate of the richest-1% was identical for all subjects, controlling for $(w_{i,t-1}^{feed} - w_{i,t-1}^{prior})$ makes it unnecessary to also control for $w_{i,t-1}^{prior}$. Also, we do not include any additional controls (i.e., X_i is empty) in the baseline specification, but we add them in alternative specifications for robustness checks.

of a given piece of information may be smaller when delivered alongside other messages.

6.4 Effect of Expected Savings

The 2SLS estimates are reported in Table 2. Panel A reports the total effects, while Panel B shows the partial effects. The specifications are identical across columns, except that each column uses a different dependent variable. In column (1), the dependent variable is our primary outcome of interest: an indicator for whether the subject filed an appeal in 2022, based on administrative data.

We begin by describing the effects of the belief about expected savings, as it serves as a natural benchmark for interpreting the remaining results. According to Panel A, which corresponds to total effects, the coefficient in column (1) on expected savings is positive (1.302), statistically significant (p-value = 0.022), and economically meaningful. This coefficient implies that a $100 increase in expected savings increases the probability of appealing by 1.30 pp—equivalent to 2.8% of the baseline appeal rate. This result is consistent with basic economic theory: individuals respond to financial incentives. Moreover, this estimate is broadly consistent with evidence from Nathan et al. (2025), who provides related quasi-experimental evidence. Using a regression kink design, they show that a $100 increase in expected tax savings raises the probability of appealing by 2.14 pp.[51] We would not expect the magnitudes to match exactly across the two studies, given differences in time periods, samples, and that each of the estimates are subject to statistical uncertainty. Nonetheless, the fact that the two estimates are directionally consistent and of the same order of magnitude (1.30 pp vs. 2.14 pp) provides reassuring evidence of the validity of our findings.

We conduct several robustness checks for this estimate. The first is a falsification test that leverages the timing of the information intervention in an event-study framework. Specifically, we re-estimate the same regression but use appeals during each of the five previous years as the dependent variable. Intuitively, because the information was provided in 2022, it could not have influenced appeal decisions in any year prior to 2022. We therefore expect the coefficients from each falsification test to be close to zero and statistically insignificant. These event-study results are shown in the left panels of Figure 7. Panel A corresponds to the total effect of expected savings. The rightmost coefficient represents the post-treatment effect—the same coefficient of 1.302 from Table 2 discussed above. The remaining coefficients correspond to each of the pre-treatment years. As expected, the coefficients are close to zero and statistically insignificant for all pre-treatment years, while the coefficient for the post-treatment year is positive and statistically significant.

[51] This estimate is reported on page 285 of Nathan et al. (2025).

The right panels of Figure 7 present another robustness check using binned scatterplots. The 2SLS models described above assume that the effect of beliefs on behavior is linear. The belief-updating results discussed earlier suggest that the linear assumption is an excellent approximation for the first stage. However, the reduced-form relationship may deviate from linearity. For example, individuals may react asymmetrically to information—they might be more likely to appeal when expected savings increase but fail to change their behavior when expected savings decrease. There could also be other non-linearities or, more concerning, the results could be driven by outliers. To examine this possibility more closely, the right panels of Figure 7 shows the intention-to-treat effects of providing information using binned scatterplots. Specifically, Panel B corresponds to the effects of information shocks about expected savings.[52] For brevity, this panel pools subjects with and without binding homestead caps, but Figure A.5 reports the results for the two groups separately. The results suggest that the findings are not driven by outliers or non-linearities.

While Panel A of Table 2 reports the total effects, Panel B reports the partial effects. Given that learning is highly compartmentalized, we would expect the total effects and partial effects to be similar to each other. Indeed, like the total effect, the partial effect is positive and statistically significant (p-value = 0.005). There are some differences, though. First, the partial effect is estimated more precisely than the corresponding total effect—this is largely mechanical, as the total effect is estimated using only two of the eight treatment groups, whereas the partial effect uses variation from all eight treatment groups. Second, in terms of magnitude, the partial effect (0.583) is roughly half the size of the total effect (1.302), but this difference must be interpreted cautiously, as it is statistically insignificant (p-value=0.235).[53] With that caveat in mind, one possible interpretation is information saturation: when individuals receive multiple pieces of information simultaneously, their attention may be divided, leading to a weaker response to any single piece.

We conduct a series of additional robustness checks. A common concern in 2SLS estimation is the possibility of weak-instrument bias (Stock et al., 2002; Andrews and Stock, 2005). However, the substantial belief updating documented earlier suggests that instrument strength is unlikely to be a problem in our setting. To formally assess instrument relevance, Table 2 reports the Kleibergen-Paap *rk* Wald F-statistics for each 2SLS regression, a stan-

[52] More precisely, this panel is similar to Panel B of Figure 6, except that the y-axis reports the probability of appealing in 2022 rather than the posterior belief about expected savings.

[53] We test the equality of two coefficients estimated in separate regressions using the same sample. Since the estimates come from different specifications, their sampling covariance is not observed; for tractability, we assume it is zero.

dard diagnostic for weak instruments.[54] In all cases, these statistics exceed the conventional rule of thumb from Staiger and Stock (1997), which suggests that if the F-statistic is above 10, weak identification is not a concern.

Additional robustness checks are presented in Appendix A.8, which reports sensitivity analyses of the baseline specification. These include adding extra control variables, using alternative definitions of outliers, and excluding subjects who did not pass the attention check. The results remain qualitatively consistent and quantitatively similar across all alternative specifications.

Column (2) of Table 2 is identical to column (1), except that the dependent variable is the ex-ante intention to file an appeal, as measured in the survey data. More specifically, the dependent variable equals 1 for individuals who reported they would "almost surely" file an appeal in 2022 and 0 otherwise.[55] The correlation between the intention to appeal and actual appeals is far from perfect: in the control group, the correlation coefficient is 0.32. In other words, many individuals who say they will appeal ultimately do not, and conversely, some who say they will not appeal end up doing it. Given this low correlation, we would not expect the effects to be mechanically the same for actual appeals (column (1)) and intentions to appeal (column (2)). We do not find any significant positive effects on the intention to appeal. This suggests that while the information on expected savings ultimately influenced the decision to file an appeal, the effect may not have been immediate—individuals may have reflected on the information for days or weeks before it affected their final decision. An alternative explanation is social desirability bias: when presented with information suggesting that filing an appeal is less profitable than expected, individuals might be reluctant to state that they no longer intend to appeal, as doing so could signal that they file appeals for selfish reasons rather than on principle.[56]

Column (3) of Table 2 shows the effects on agent appeals, an outcome where we do not expect to find any effects.[57] As expected, we find no significant effect of providing owners

[54] Since the rule of thumb is derived under the assumption of homoskedastic errors, it would be technically appropriate to use the Cragg-Donald statistic, which is also derived under homoskedastic errors. However, we follow the recommendation of Baum et al. (2007): "the use of the rk Wald statistic, as the robust analog of the Cragg-Donald statistic, is a sensible choice and clearly superior to the use of the latter in the presence of heteroskedasticity."

[55] The number of observations is slightly smaller in column (2) than in column (1) because a small share of subjects left the survey after the elicitation of posterior beliefs but before answering the question on intent to appeal.

[56] For related evidence on social desirability bias in the context of tax preferences, see Ajzenman et al. (2024).

[57] We do not expect to find any effects on this outcome because we offer subjects the same services agents provide for free as an incentive for participating in the experiment.

with information about expected savings on their likelihood of appealing through agents.[58] In turn, column (4) of Table 2 shows the effects on successful appeals. In other words, this outcome indicates whether the owner appealed directly *and* such appeal resulted in a reduction in the proposed value, and 0 otherwise. This measure allows us to assess whether the marginal appeals induced by the treatment were consequential. In both Panels A and B, the coefficient estimates in column (4) are similar to those in column (1) in terms of sign, magnitude, and statistical significance. Thus, the data suggest that the marginal appeals induced or discouraged by the information on expected tax savings were largely consequential.

6.5 Effect of Perceived Appeals by the Richest-1%

We now turn to the effects of the belief about the richest-1% appeal rate. Evidence from Section 4 suggests that individuals are less tolerant of tax avoidance by the wealthiest households. A natural starting point, therefore, is to examine whether learning that the wealthy appeal more frequently alters individuals' fairness perceptions. The results are presented in column (5) of Table 2. The dependent variable is the answer to the survey question on the fairness of the gap in appeal rates between richest-1% and comparable households, on a scale from 0 (very unfair) to 10 (very fair). The average fairness score is 4.8.[59] For easier interpretation, column (5) uses a standardized outcome.

The total effects, shown in Panel A of Table 2, are imprecisely estimated and statistically insignificant.[60] It is important to keep in mind that this is a subjective outcome, which is typically noisy and therefore produces less precise estimates. However, the partial effects, shown in Panel B, provide suggestive evidence that a larger perceived gap in appeal rates between the richest-1% and comparable homes reduces perceived fairness. More precisely, each 1 pp increase in the perceived appeal rate of the richest-1% *reduces* the fairness score by 0.005 standard deviations (p-value=0.003). Likewise, each 1 pp increase in the perceived appeal rate of comparable homes *increases* the fairness score by 0.006 standard deviations (p-value=0.008). By contrast, perceptions of expected savings have an effect on perceived fairness that is close to zero and statistically insignificant. This experimental result aligns with the evidence in Section 4, which shows that individuals are less tolerant of tax avoidance by the wealthiest households. However, as we demonstrate below, despite the effects on this

[58] Consistent with this evidence, prior experiments in the same context have found that information provided to owners affects owner appeals but has no impact on agent appeals (Nathan et al., 2025; Giaccobasso et al., 2025; Nathan et al., 2023).

[59] For more details, see Figure A.7.

[60] The only exception is the coefficient on expected savings, which is positive and borderline significant in Panel A. However, since the corresponding coefficient in Panel B is near zero, precisely estimated, and statistically insignificant, we view the Panel A result as likely spurious.

survey outcome, there is no significant impact on the actual decision to file a tax appeal.

In column (1) of Table 2, the dependent variable is whether the individual filed a tax appeal. Panel A shows that the total effect of the richest-1% appeal rate is close to zero (-0.034), precisely estimated, and statistically insignificant (p-value = 0.810). This implies that a 1 pp increase in the appeal rate of the richest-1% would, if anything, *reduce* the subject's own probability of filing an appeal by 0.034 pp. This null effect runs counter to the contagion hypothesis. However, before delving further into the magnitude of the coefficient, we first examine its robustness.

Panel C of Figure 7 presents the event-study analysis for the total effect of the richest-1% appeal rate. All pre-treatment coefficients are close to zero and statistically insignificant, which supports the validity of the post-treatment coefficient. Panel D of Figure 7 shows the intention-to-treat effects in binned scatterplot form. The null finding does not appear to be driven by non-linearities or outliers. The results of column (1) of Table 2 remain very similar across alternative specifications (see Appendix A.8) and are also null for the intention to appeal, agent appeals, and successful appeals (columns (2)–(4) of Table 2).

Panel B of Table 2 reports the partial effects. Because learning is highly compartmentalized, we expect small or no differences between the partial and total effects. Indeed, the coefficient for the partial effect remains close to zero (0.126) and statistically insignificant. On the positive side, the partial effect is more precisely estimated and, for this reason, may be preferred to the total effect. On the other hand, the event-study analysis for the partial effect shows some minor anomalies—see Appendix A.8 for details—and for this reason we prefer to focus on the total effect as the headline estimate.

Now that we have established the robustness of this finding, we can return to the discussion of its economic magnitude. The total effect of the wealthy appeal rate (-0.034, from Panel A, column (1) of Table 2) is not even directionally consistent with the hypothesis of trickle-down effects, which predicts positive effects. However, the fact that the coefficient is close to zero and statistically insignificant does not necessarily mean it is a precise null. To assess this, we examine the upper bound of the 90% confidence interval, which allows us to rule out effects larger than 0.199 at the 90% confidence level. This would be a modest effect: a 1 percentage point increase in richest-1% appeals would increase the probability of filing an appeal by at most 0.199 pp.[61]

One natural benchmark for this bound is the effect of expected savings. According to this upper bound, the effect of a 1 percentage point increase in richest-1% appeals would be, at most, equivalent to the effect of an \$15 increase in expected savings ($= \$100 \cdot \frac{0.199}{1.302}$). Another

[61] The conclusions would be similar if we used the partial effect from Panel B, where the upper bound of the 90% confidence interval is 0.268.

natural benchmark is Giaccobasso et al. (2025), who study the effects of a belief which, like the richest-1% appeal rate, is also expressed as a share: the percent of property taxes that go to public schools. According to the reciprocity hypothesis, households with children—who benefit from school taxes—would be less likely to file a tax appeal when believing that a higher share of their taxes go to schools. Indeed, Giaccobasso et al. (2025) report that, for a household with children enrolled in public schools, a 1 pp increase in the perceived share of property taxes going to public schools decreases the probability of appealing by 0.409 pp. By comparison, our estimate indicates that a 1 pp increase in the perceived share of richest-1% households appealing would increase the subject's own probability of appealing by, at most, 0.199 pp.

One common concern about null effects in experiments is that participants may not be paying attention to the information provided. However, we find significant effects for the treatment on expected tax savings, indicating that participants are indeed attentive in this setting. Another potential concern is that tax appeal decisions may be influenced by information on pecuniary incentives but not by other types of information, such as those appealing to tax morale. Yet, evidence from other information-provision experiments conducted in the same context using similar methodologies suggests otherwise. For example, Giaccobasso et al. (2025) show that informing households about how their tax dollars are spent—specifically, whether they fund services that benefit them, such as local public schools—significantly affects subsequent appeal rates. Similarly, Nathan et al. (2023) find that providing information about the tax rate paid by the average household in the county influences both perceived fairness and the likelihood of appealing. Taken together with these studies, our evidence suggests that it is the specific information about the appeal rates of the richest-1% that has no impact on filing decisions.

In any case, we should not draw conclusions based on a single coefficient. We use additional approaches to estimate trickle-down effects, relying on alternative dependent variables that capture the hypothesized causal mechanisms more directly. We also examine horizontal comparisons, since if vertical contagion existed, one might expect horizontal contagion as well. We present these additional results next, all of which reinforce the conclusion that perceptions about the appeal rate of the richest-1% have no effect on the decision to file a tax appeal.

6.6 Effect of Perceived Appeals by Comparable Homes

We next turn to the effects of the belief about the appeal rate of comparable homes. The two underlying channels—observational learning and social norms—predict effects in the same direction as those for vertical comparisons. In fact, there are reasons to expect horizontal

comparisons to have even stronger effects than vertical comparisons.[62] People tend to compare themselves to their peers more than to other groups (Festinger, 1954; Clark and Senik, 2010), making them more likely to form social norms based on the behavior of comparable homes. Evidence from the context of salary comparisons illustrates this dynamic: employees tend to feel jealous when their colleagues earns a bit more than they do, but do not feel jealous when their boss earns far more (Cullen and Perez-Truglia, 2022). Likewise, one may expect observational learning to be stronger when the observed behavior comes from households with similarly valued homes, as their tax appeals are more likely to reflect the individual's own costs and benefits. Experts in our forecast survey echoed this logic, with most predicting that beliefs about comparable homes would have a greater effect on appeal decisions than beliefs about the richest-1%—for details, see Appendix B.

In column (1) of Table 2, the total effect (shown in Panel A) is close to zero (-0.054) and statistically insignificant (p-value = 0.730). This implies that a 1 pp increase in the appeal rate of comparable homes would, if anything, *reduce* the subject's own probability of filing an appeal by 0.054 pp. Although small, the negative coefficient runs counter to the contagion hypothesis. Furthermore, all the robustness checks presented above for the richest-1% apply similarly to the coefficient on comparable homes. The total effect for the belief about comparable homes (-0.054) is qualitatively and quantitatively similar to the corresponding coefficient for the belief about the richest-1% (-0.034). The event-study analysis shows no effect on pre-treatment outcomes (Panel E of Figure 7). The binned scatterplot indicates that the null effect does not seem to be driven by non-linearities or outliers (Panel F of Figure 7). The results remain very similar across alternative specifications (see Appendix A.8). There are also null effects for the intention to appeal, agent appeals, and successful appeals (columns (2)–(4) of Table 2). Given the compartmentalized learning, the partial and total effects should be similar. Indeed, like the total effect, the partial effect is also close to zero and statistically insignificant (Panel B of Table 2). In sum, the fact that we do not observe any effects from beliefs about comparable homes reinforces the conclusion that beliefs about the richest-1% also had no impact.

6.7 Direct Test of the Observational Learning Channel

We use posterior beliefs about households' own expected tax savings to directly test the observational learning channel. According to this channel, when individuals learn that the

[62] Of course, one could also argue the opposite. For example, if the wealthy are perceived as financially savvy, their decision to file appeals could amplify the observational learning response, signaling more strongly that appealing is a smart strategy. Likewise, learning that the wealthy appeal at high rates might trigger stronger moral judgment or even a sense of retaliation.

richest-1% appeal more often, they should infer that appealing must be financially advantageous and therefore raise their own expected savings. The same logic applies to horizontal comparisons: when individuals learn that comparable homes appeal more frequently, they should become more optimistic about their own expected savings.

The evidence is presented in Figure 8. This figure follows the same specification described earlier in Section 5.2, which we used to measure learning. However, while the previous section examined whether information about a belief led to updates about that same belief, Figure 8 measures *cross-learning*—that is, whether information about one belief led to updates in other beliefs. Specifically, Panel A compares two treatment groups: individuals who did not receive any information versus individuals who received only information about the appeal rates of the richest-1%. The x-axis represents the belief gap—whether the household underestimated or overestimated the appeal rates of the richest-1%. The y-axis, however, is *not* the belief update about the appeal rate of the richest-1% but rather the belief update about the household's own expected tax savings. The difference in slopes between treatment and control groups is close to zero (0.018) and statistically insignificant. This means that finding out that the richest-1% appeal rate is 1 pp higher than you thought increases the household's own expected savings by merely \$1.8. Moreover, this effect is estimated very precisely around zero. Using the upper bound of the corresponding 90% confidence interval, we can rule out, with 90% confidence, that when an individual finds out that the appeal rate of the richest-1% is 1 pp higher than expected, they increase their expected tax savings by at most \$8.3. This reflects a precisely bounded null effect.

Panel B of Figure 8 mirrors Panel A but focuses on horizontal comparisons rather than vertical comparisons. Even for horizontal comparisons, the degree of cross-learning is statistically insignificant and tightly estimated around zero. For example, based on the difference in slopes from Panel B, when finding out that the comparable appeal rate is 1 pp higher than expected, on average households increase their own expected savings by just \$4.7. And according to the corresponding confidence interval, we can rule out, with 90% confidence, effects above \$11.13. This is again a precisely estimated null. So, the fact that individuals do not *even* exhibit observational learning about comparable homes further reinforces the finding of no observational learning about the richest-1%.

6.8 Direct Test of the Social Norms Channel

Next, we explore the social norms channel. In social psychology (Deutchman et al., 2025), it is believed that individuals infer social approval (*injunctive* norms) from what they perceive others are doing (*descriptive* norms). According to this social norms channel, when individuals observe that a high share of the richest-1% are appealing (the descriptive norm), they may

feel it is more acceptable to file an appeal (the injunctive norm). To test this hypothesis, we use a question from the field experiment that directly asks individuals about their injunctive norms: i.e., whether they think tax appeals are acceptable. In the control group, 83.01% said tax appeals are always justifiable, 16.75% said they are sometimes justifiable, and only 0.24% deemed it never justifiable. While we did not use the same question in the Attitudes Survey and the field experiment, this relatively high acceptability of tax appeals is broadly consistent with the complementary evidence presented in Section 4.1 above.

The results are presented in column (6) of Table 2, which is identical to column (1) except that the dependent variable is now an indicator for whether the individual stated that tax appeals are always justifiable.[63] The total effect of the belief about the appeal rates of the richest-1% is close to zero (0.111) and statistically insignificant. In other words, when households learn that the richest-1% are appealing more often, it does not change whether they view tax appeals as justifiable. When looking at the partial effect of this belief, the coefficient is also close to zero (0.090) and statistically insignificant. For horizontal comparisons, however, the results are mixed, as the partial effect is close to zero (0.053) and statistically insignificant, but the total effect is larger (0.262) and statistically significant (p-value=0.021). So, for horizontal comparisons, learning that peers appeal more often may have made appealing seem more acceptable, but the effect was likely too weak to significantly influence whether the subject ultimately chose to appeal.

7 Conclusions

Tax avoidance is widespread among the wealthiest individuals, and the public is increasingly aware of these practices. This growing visibility raises concerns that exposure to elite avoidance could erode tax compliance more broadly—either by highlighting replicable strategies or by normalizing avoidance as socially acceptable. We provide a test of this hypothesis using novel data. First, we use survey and willingness-to-pay data to show that people are substantially less accepting of tax avoidance when incurred by the most wealthy. Second, we use large-scale, pre-registered field experiment in Dallas County. When households find out that the wealthiest are more likely to file and appeal, it reduces their perceptions of fairness. However, updating beliefs about how common tax avoidance is among the wealthy does not increase the likelihood of appealing, nor does it shift beliefs about the financial or moral acceptability of appealing. Consistently, learning about the frequency of tax appeals among comparable households does not have effects either. By comparison, individuals who

[63] While the question has a 3-point scale, given that virtually nobody chose the lowest point, we collapsed the variable into a simple indicator outcome, for ease of interpretation.

learn that appeals can yield substantial tax savings are significantly more likely to file. Taken together, these findings suggest that although the public disapproves of tax avoidance by the wealthy, their behavior does not trickle-down to the broader population.

A common consideration in any empirical study is the external validity of the findings. On the one hand, these results may be generalizable to other contexts. For example, to the extent that property tax appeals operate similarly in other Texas counties, as well as in counties outside of Texas (Dobay et al., 2019; World Bank, 2019; Nathan et al., 2025), the findings are likely to apply more broadly.[64] On the other hand, there are some limitations worth noting. First, due to our recruitment strategy, our subject pool overrepresents households that were more inclined to file a property tax appeal. As a result, the effect sizes might have been smaller had we used a more representative sample. However, given that we document a null effect, this potential bias only strengthens our conclusion.

Another important consideration is that our study focuses on one specific form of tax avoidance—property tax appeals. The two mechanisms we study may be weaker or stronger for other types of tax avoidance. Indeed, our context arguably provides ideal conditions for observational learning. In other settings, the strategies used by the wealthiest individuals are often very different from those suitable for most Americans. For example, the buy-borrow-die strategy is sensible for individuals with vast stock wealth but is irrelevant for the vast majority of households. By contrast, the tax appeal process is identical for every household, regardless of wealth. Filing an appeal requires completing the same simple form, whether the home is worth $50,000 or $50 million. Since this context provides ideal conditions for observational learning, the absence of any evidence here suggests it is even less likely to occur in other settings beyond tax appeals. On the flip side, it could be argued that our context is not particularly favorable for the social norms channel. Although not everyone approves of tax appeals, baseline acceptance is already quite high, leaving limited room for acceptance to increase further. Thus, the social norms channel could play a larger role in other, more complex or stigmatized forms of tax avoidance or tax evasion. In any case, further research is needed to assess whether and how these effects generalize. In the terminology of List (2020), we view our results as a Wave 1 insight—establishing initial causal relationships and providing first tests of theory.

Our findings have implications for tax transparency and communication strategies around tax policy. Recent efforts by researchers (e.g., Zucman, 2024), government agencies (Leiserson and Yagan, 2021), and non-governmental organizations (ProPublica, 2023) publicize

[64] There are exceptions, however—for instance, property tax appeals are relatively rare in California due to Proposition 13. Since a property's assessed value is capped and only reassessed upon sale or major improvements, most long-time property owners pay taxes well below market value—giving them no incentive to file appeals.

the effective tax rates of the wealthy and the strategies they use to avoid taxes. A potential concern is that such publicity could unintentionally lower tax compliance among the broader population. Similarly, there is a concern that efforts to increase tax transparency—such as naming and shaming tax delinquents—could undermine tax compliance more broadly if they lead the public to believe that tax avoidance and evasion are more widespread than previously assumed (see e.g., Perez-Truglia and Troiano, 2018). While further research is needed, our evidence suggests that the unintended negative consequences of this increased transparency may be less severe than initially feared.

References

Abadie, A., J. Gu, and S. Shen (2024). Instrumental Variable Estimation with First-Stage Heterogeneity. *Journal of Econometrics 240*(2), 105425.

Ajzenman, N., G. Cruces, R. Perez-Truglia, D. Tortarolo, and G. Vazquez-Bare (2024). From Flat to Fair? The Effects of a Progressive Tax Reform. *NBER Working Paper No. 33286*.

Alstadsæter, A., N. Johannesen, S. L. G. Herry, and G. Zucman (2022). Tax Evasion and Tax Avoidance. *Journal of Public Economics 206*, 104587.

Alstadsæter, A., N. Johannesen, and G. Zucman (2019). Tax Evasion and Inequality. *American Economic Review 109*(6), 2073–2103.

Alstadsæter, A., W. Kopczuk, and K. Telle (2019). Social networks and tax avoidance: evidence from a well-defined norwegian tax shelter. *International Tax and Public Finance 26*(6), 1291–1328.

Andrews, D. W. and J. H. Stock (2005). *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*. Cambridge University Press.

Battaglini, M., L. Guiso, C. Lacava, and E. Patacchini (2020). Tax Professionals and Tax Evasion. *NBER Working Paper No. 25745*.

Baum, C. F., M. E. Schaffer, and S. Stillman (2007). Enhanced routines for instrumental variables/generalized method of moments estimation and testing. *The Stata Journal 7*(4), 465–506.

Bergolo, M., R. Ceni, G. Cruces, M. Giaccobasso, and R. Perez-Truglia (2020). What Makes a Tax Evader? *NBER Working Paper No. 28235*.

Best, M. C., L. Caloi, F. Gerard, E. P. Kresch, J. Naritomi, and L. Zoratto (2025). Greener on the other side: Inequity and tax compliance. *NBER Working Paper No. 34062*.

Boning, W. C., J. Guyton, R. Hodge, and J. Slemrod (2020). Heard it through the grapevine: The direct and network effects of a tax enforcement field experiment on firms. *Journal of Public Economics 190*, 104261.

Bottan, N. L. and R. Perez-Truglia (2025). Betting on the house: Subjective expectations and

market choices. *American Economic Journal: Applied Economics 17*(1), 459–500.

Brown, J., S. Jilke, D. Schwegman, M. Toma, and M. C. Turner (2025). Improving Tax Compliance among the Clients of Tax Preparers. *SSRN Working Paper No. 5400632*.

Bursztyn, L. and R. Jensen (2017). Social image and economic behavior in the field: Identifying, understanding, and shaping social pressure. *Annual Review of Economics 9*, 131–153.

Cai, H., Y. Chen, and H. Fang (2009). Observational Learning: Evidence from a Randomized Natural Field Experiment. *American Economic Review 99*(3), 864–82.

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

Çelebi, C., C. Exley, S. Harrs, H. Kivimäki, M. Serra-Garcia, and J. Yusof (2025). Mission possible: Data quality in online surveys.

Chetty, R., J. N. Friedman, and E. Saez (2013). Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings. *The American Economic Review 103*(7), 2683–2721.

Clark, A. E. and C. Senik (2010). Who Compares to Whom? The Anatomy of Income Comparisons in Europe. *The Economic Journal 120*(544), 573–594.

CNN (2015). Trump Unveils Tax Plan, Press Conference Transcript, aired September 28, 2015.

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

Cruces, G., D. Tortarolo, and G. Vazquez-Bare (2022). Design of Two-Stage Experiments with an Application to Spillovers in Tax Compliance. *Working Paper*.

Cullen, J. B., N. Turner, and E. Washington (2021). Political Alignment, Attitudes toward Government, and Tax Evasion. *American Economic Journal: Economic Policy 13*(3), 135–66.

Cullen, Z. and R. Perez-Truglia (2022). How much does your boss make? the effects of salary comparisons. *Journal of Political Economy 130*(3), 766–822.

DellaVigna, S., D. Pope, and E. Vivalt (2019). Predict science to improve science. *Science 366*(6464), 428–429.

Deutchman, P., G. Kraft-Todd, L. Young, and K. McAuliffe (2025). People update their injunctive norm and moral beliefs after receiving descriptive norm information. *Journal of Personality and Social Psychology 128*(1), 1–18.

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

Drago, F., F. Mengel, and C. Traxler (2020). Compliance Behavior in Networks: Evidence from a Field Experiment. *American Economic Journal: Applied Economics 12*(2), 96–133.

Edwards, A. L. (1957). *The Social Desirability Variable in Personality Assessment and Research.* New York: Dryden Press.

Feldman, N. E., P. Katuščák, and L. Kawano (2016). Taxpayer confusion: Evidence from the child tax credit. *American Economic Review 106*(3), 807–835.

Feldman, N. E. and J. Slemrod (2007). Estimating tax noncompliance with evidence from unaudited tax returns. *The Economic Journal 117*(518), 327–352.

Festinger, L. (1954). A theory of social comparison processes. *Human Relations 7*(2), 117–140.

Fuster, A., R. Perez-Truglia, M. Wiederholt, and B. Zafar (2022). Expectations with Endogenous Information Acquisition: An Experimental Investigation. *Review of Economics and Statistics 104*(5), 1059–1078.

Giaccobasso, M., B. Nathan, R. Perez-Truglia, and A. Zentner (2025). Where do my tax dollars go? Tax morale effects of perceived government spending. *American Economic Journal: Applied Economics 17*(4), 223–59.

Holz, J., R. Jiménez-Durán, and E. Laguna-Müggenburg (2024). Estimating the distaste for price gouging with incentivized consumer reports. *American Economic Journal: Applied Economics 16*(1), 33–59.

Kleven, H. J., C. Landais, and E. Saez (2013). Taxation and International Migration of Superstars: Evidence from the European Football Market. *American Economic Review 103*(5), 1892–1924.

Landier, A. and G. Plantin (2017). Taxing the Rich. *The Review of Economic Studies 84*(3), 1186–1209.

Leiserson, G. and D. Yagan (2021). What is the Average Federal Individual Income Tax Rate on the Wealthiest Americans? *Council of Economic Advisers' Blog, posted on September 23, 2021.*.

Levitt, S. D. and J. A. List (2007). What do laboratory experiments measuring social preferences reveal about the real world? *Journal of Economic perspectives 21*(2), 153–174.

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

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

Nathan, B., R. Perez-Truglia, and A. Zentner (2023). Paying Your Fair Share: Perceived Fairness and Tax Compliance. *NBER Working Paper No. 32588*.

Nathan, B., R. Perez-Truglia, and A. Zentner (2025). My Taxes Are Too Darn High: Why Do Households Protest Their Taxes? *American Economic Journal: Economic Policy 17*(1), 273–310.

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

Perez-Truglia, R. and J. Yusof (2024). Billionaire Superstar: Public Image and Demand for Taxa-

tion. *NBER Working Paper No. 32712*.

ProPublica (2023). America's Highest Incomes and Taxes, Revealed. https://projects.propublica.org/americas-highest-incomes-and-taxes-revealed/. Accessed: 2024-09-18.

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

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

Slemrod, J. and S. Yitzhaki (2002). *Chapter 22: Tax Avoidance, Evasion, and Administration, in Handbook of Public Economics, edited by: Auerbach, A.J. and Feldstein, M.* Elsevier.

Smith, M., D. Yagan, O. Zidar, and E. Zwick (2019). Capitalists in the Twenty-First Century. *The Quarterly Journal of Economics 134*(4), 1675–1745.

Staiger, D. and J. H. Stock (1997). Instrumental variables regression with weak instruments. *Econometrica 65*(3), 557–586.

Stock, J. H., J. H. Wright, and M. Yogo (2002). A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business & Economic Statistics 20*(4), 518–529.

Taylor, J. (2016). Clinton And Trump Clash In Tense First Presidential Debate. *NPR Politics*. Published Sept. 26, 2016.

Tyger, A. and S. Eastman (2019). The Regressivity of Deductions. *Tax Foundation Blog*. Accessed: 2024-10-25.

U.S. Census Bureau (2021). Population, Dallas County, Texas. *https://www.census.gov/quickfacts/fact/table/dallascountytexas/POP010220*.

Wilson, R. (2022). The Impact of Social Networks on EITC Claiming Behavior. *The Review of Economics and Statistics 104*(5), 929–945.

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

Zucman, G. (2024). Opinion: Global Billionaires Should Pay More Taxes. *The New York Times*, May 3. Accessed: 2024-09-18.

Figure 1: Sample Photos of Most and Least Expensive Homes

A. Richest-1% Homes

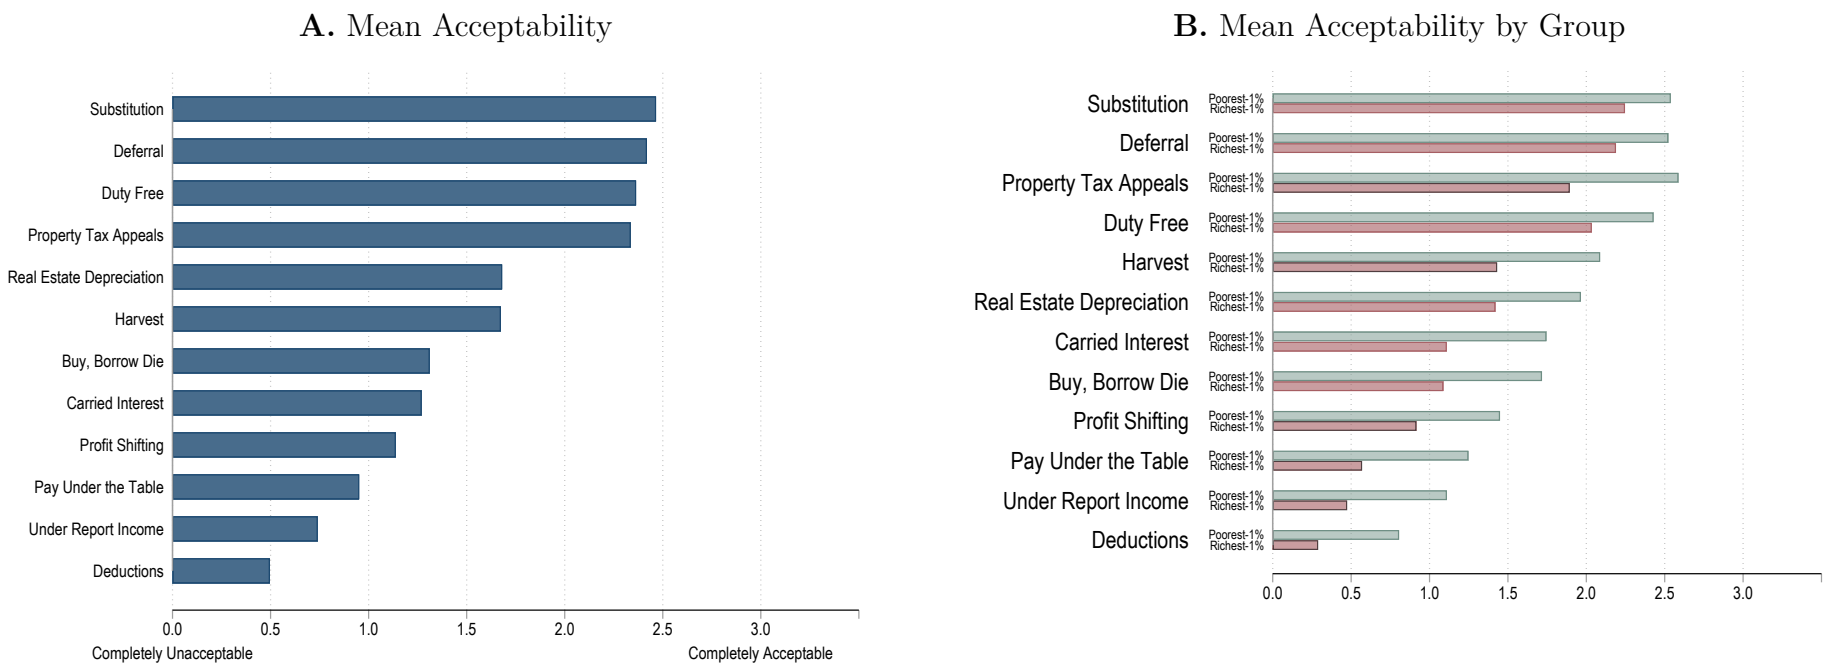


B. Poorest-1% Homes



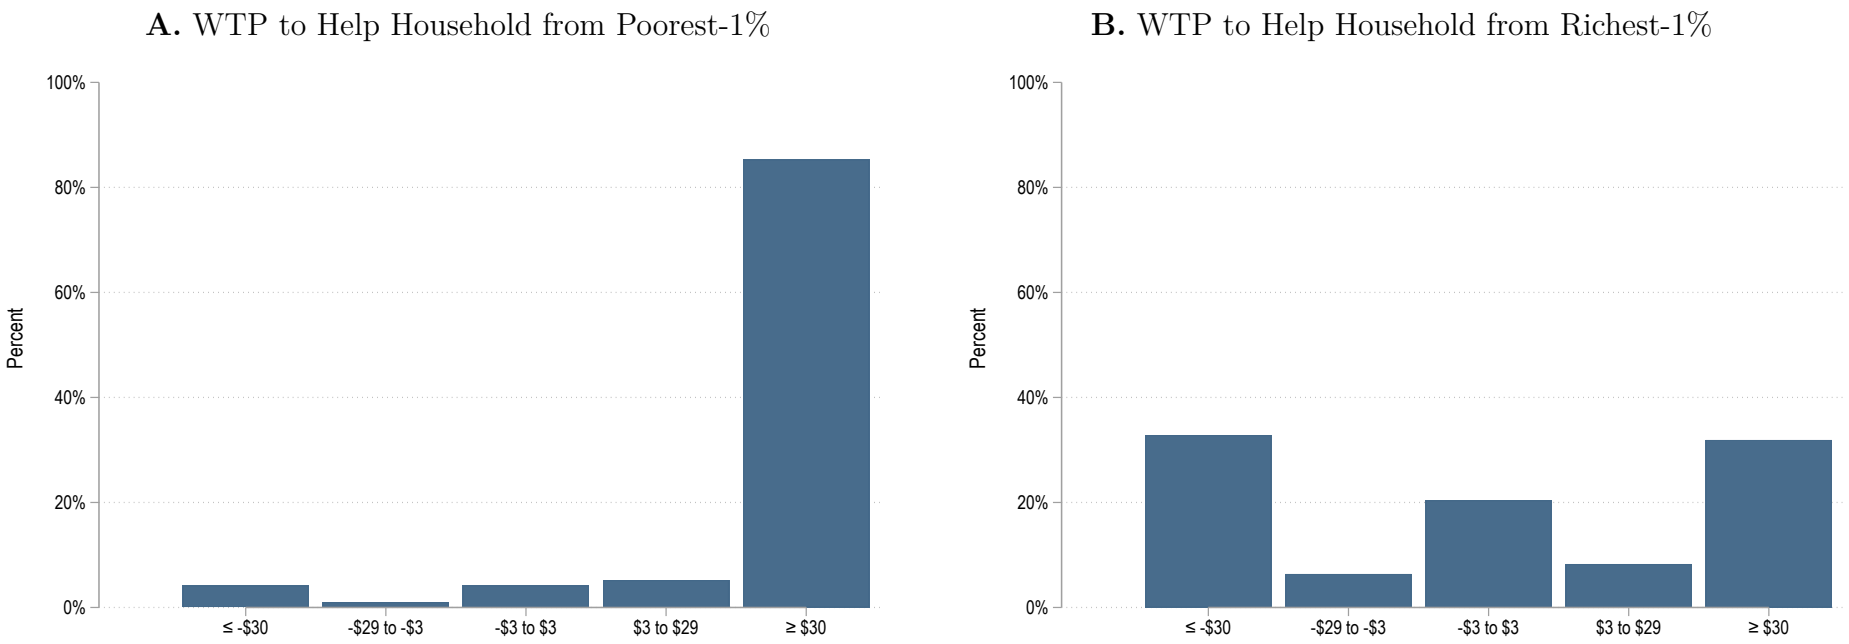
Notes: Panel A displays the images of the top 1% most expensive houses shown to the participants. Panel B displays the images of the least expensive homes shown to participants. These photos were downloaded on February 15, 2022, from a sample of homes listed for sale on http://www.redfin.com.

Figure 2: Acceptability of Tax Avoidance/Evasion Strategies



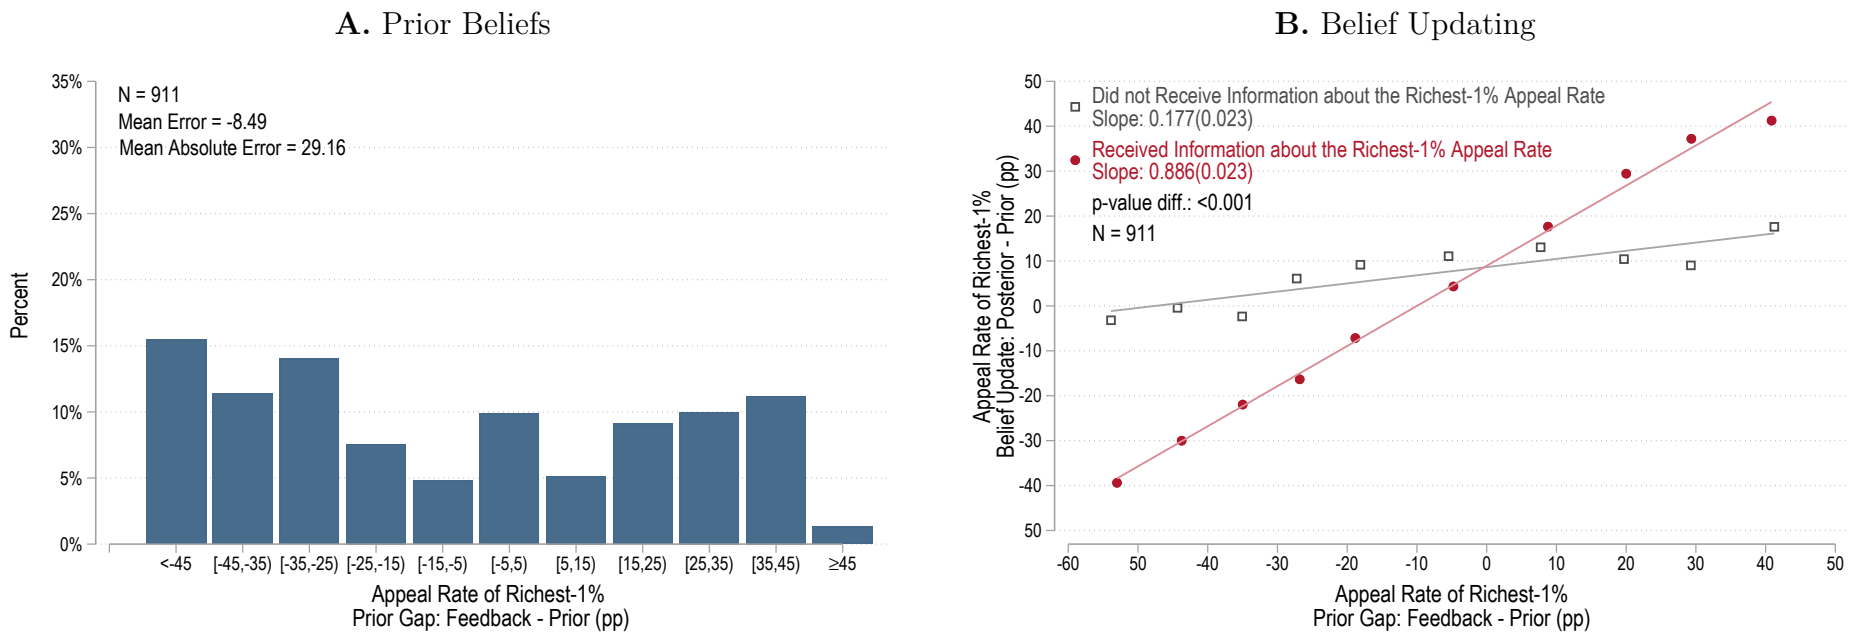
Notes: This figure uses data from the Attitudes Survey. Panel A displays responses to the question "How acceptable do you believe it is to use [strategy] appeals to lower your taxes?" coded on a 0–3 scale from completely unacceptable (0) to completely acceptable (3). Panel B displays the responses to two questions, also coded on a 0–3 scale, that distinguish between the group that is engaging in the behavior: "How acceptable do you believe it is for the poorest-1% of households to use [strategy] to lower their taxes?" and "How acceptable do you believe it is for the richest-1% of households to use [strategy] to lower their taxes?"

Figure 3: Revealed-Preference Evidence on the Acceptability of Tax Appeals



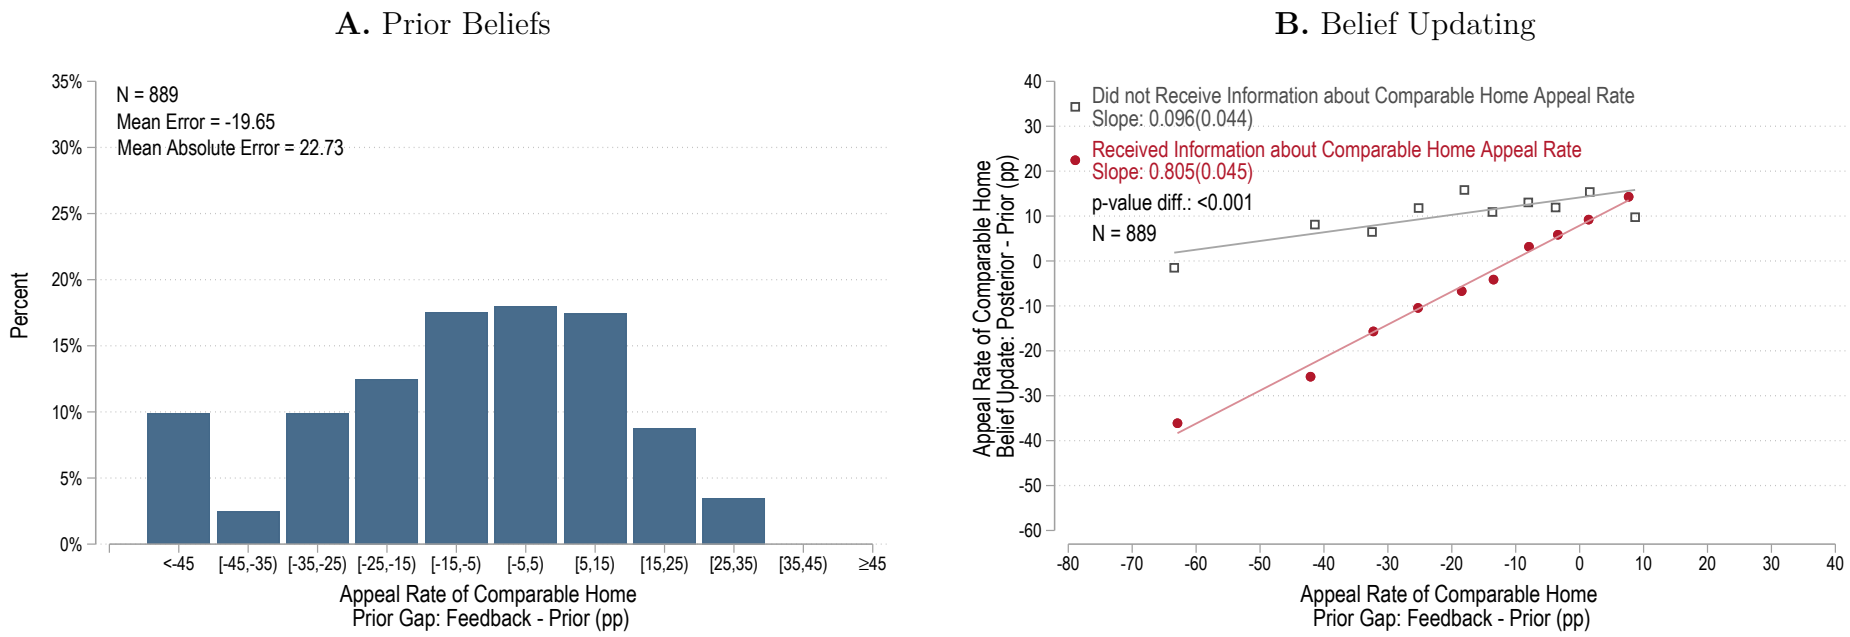
Notes: Panel A displays the distribution of the willingness to pay to help a taxpayer from the poorest-1% of the house value distribution appeal their tax bill. Panel B displays the corresponding willingness to pay to help a taxpayer from the richest-1%. Positive values indicate a willingness to pay to help someone else file a tax appeal, while negative values indicate a willingness to give up money to prevent the help from occurring. The maximum reward was \$30 in the experimental design. Therefore, the rightmost and leftmost bars are $\geq$ \$30 and $\leq -$\$30 by construction.

Figure 4: Appeal Rate of the Richest-1%: Prior Beliefs and Belief Updating



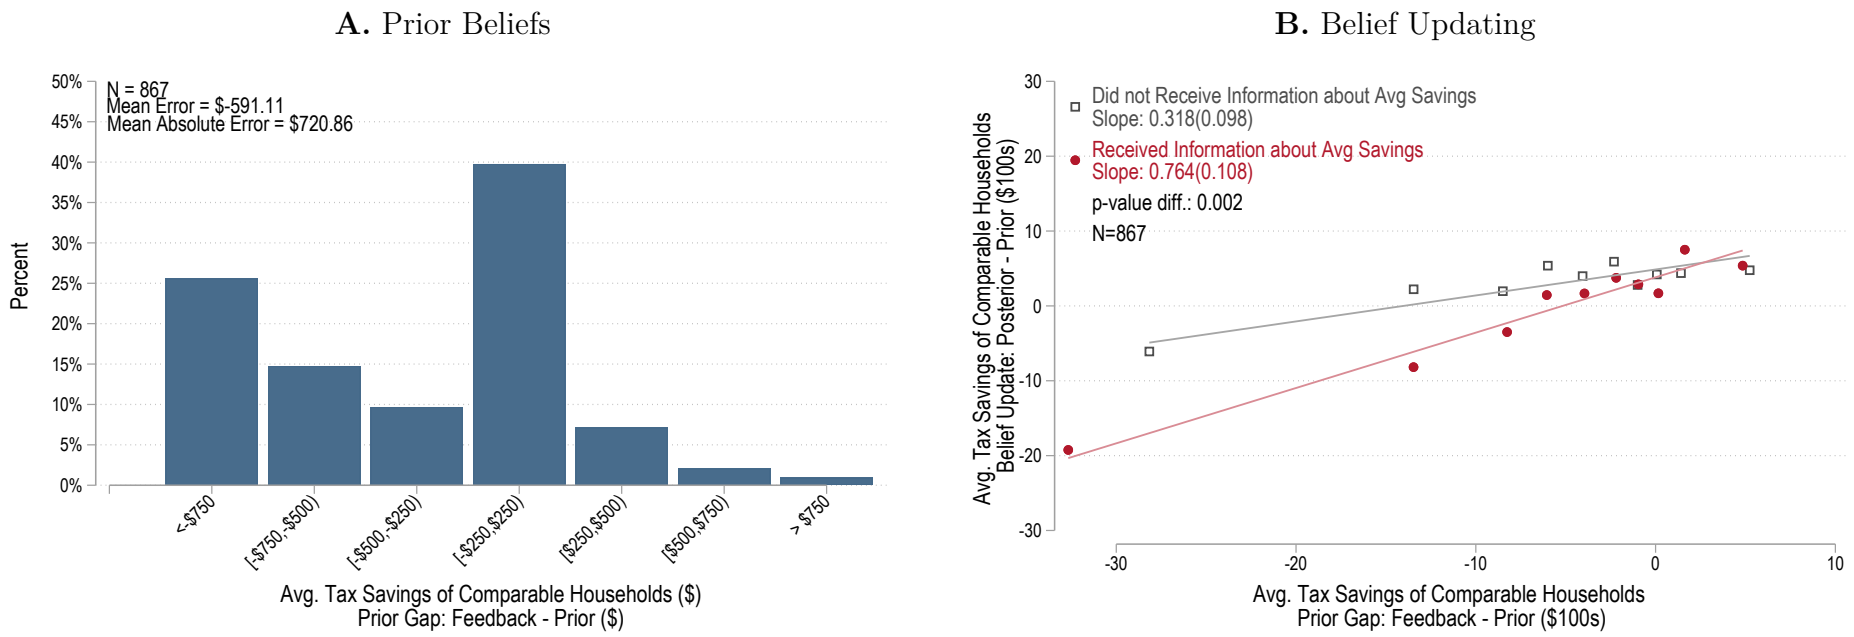
Notes: All results are based on two treatment groups: those who received information solely about the appeal rate of the richest 1% and those who did not receive any information. Panel A shows the gap in prior beliefs about the appeal rate of the richest-1%. The x-axis reports the difference between the actual richest-1% appeal rate and respondents' prior beliefs about the richest-1% appeal rate. The y-axis reports the percentage of homeowners in each bin. The upper left corner reports the total number of observations, the mean error, and the mean absolute error. Panel B shows how respondents update their beliefs using a binned scatterplot. The x-axis is the same as in Panel A. The y-axis reports the difference between posterior and prior beliefs (i.e., belief updating). Red circles (gray squares) represent the average update within each bin for the group of homeowners that were selected (were not selected) into the receiving information about the richest-1% appeal rate. Each line corresponds to the fitted values from a linear regression. The binned scatterplot include the following controls: the gap between the feedback and priors for the comparable appeal rates, the gap between the feedback and priors for tax savings interacted with a binding homestead cap and a non-binding homestead cap, and an indicator for whether the homestead cap binds. The coefficients associated with the gap variable are reported in the upper left corner, as well as their robust standard errors (in parentheses), the p-value of the difference in the slopes, and the number of observations included in the analysis.

Figure 5: Appeal Rate of Comparable Households: Prior Beliefs and Belief Updating



Notes: All results are based on two treatment groups: those who received information solely about the appeal rate of the comparable households and those who did not receive any information. Panel A shows the gap in prior beliefs about the appeal rate of comparable households. The x-axis reports the difference between the actual comparable appeal rate and respondents' prior beliefs about the comparable appeal rate. The y-axis reports the percentage of homeowners in each bin. The upper left corner reports the total number of observations, the mean error, and the mean absolute error. Panel B shows how respondents update their beliefs using a binned scatterplot. The x-axis is the same as in Panel A. The y-axis reports the difference between posterior and prior beliefs (i.e., belief updating). Red circles (gray squares) represent the average update within each bin for the group of homeowners that were selected (were not selected) into the receiving information about the comparable appeal rate. Each line corresponds to the fitted values from a linear regression. The binned scatterplot include the following controls: the gap between the feedback and priors for the richest-1% appeal rates, the gap between the feedback and priors for tax savings interacted with a binding homestead cap and a non-binding homestead cap, and an indicator for whether the homestead cap binds. The coefficients associated with the gap variable are reported in the upper left corner, as well as their robust standard errors (in parentheses), the p-value of the difference in the slopes, and the number of observations included in the analysis.

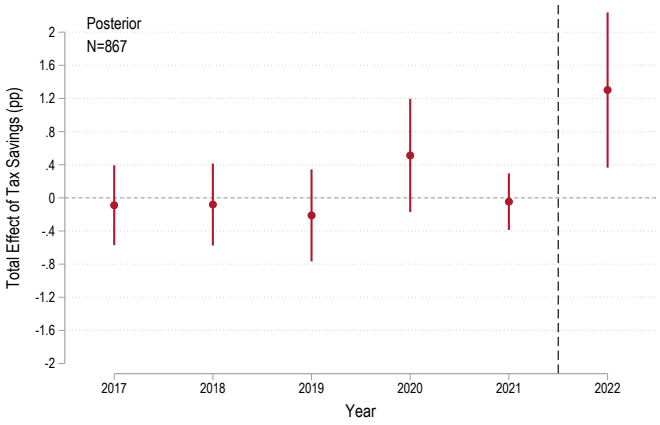
Figure 6: Expected Tax Savings: Prior Beliefs and Belief Updating



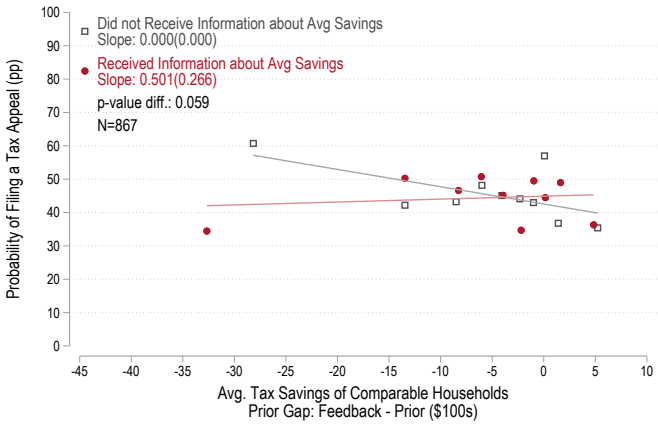
Notes: All results are based on two treatment groups: those who received information solely about the expected tax savings of comparable households and those who did not receive any information. Panel A shows the gap in prior beliefs about the expected savings. The x-axis reports the difference between the actual savings from appealing and respondents' prior beliefs about the expected savings. The y-axis reports the percentage of homeowners in each bin. The upper left corner reports the total number of observations, the mean error, and the mean absolute error. Panel B shows how respondents update their beliefs using a binned scatterplot. The x-axis is the same as in Panel A. The y-axis reports the difference between posterior and prior beliefs (i.e., belief updating). Red circles (gray squares) represent the average update within each bin for the group of homeowners that were selected (were not selected) into the receiving information about the expected savings. Each line corresponds to the fitted values from a linear regression. The binned scatterplot includes the following controls: the gap between the feedback and priors for the comparable appeal rates, and the gap between the feedback and the priors for the richest 1% appeal rates. The coefficients associated with the gap variable are reported in the upper left corner, as well as their robust standard errors (in parentheses), the p-value of the difference in the slopes, and the number of observations included in the analysis.

Figure 7: 2SLS Total Effects: Event-Study Analysis and Intention-to-Treat Binned Scatterplots

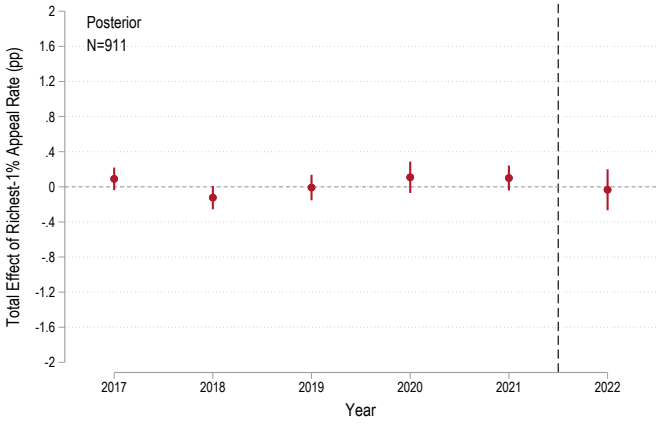
A. Expected Savings: Event-Study



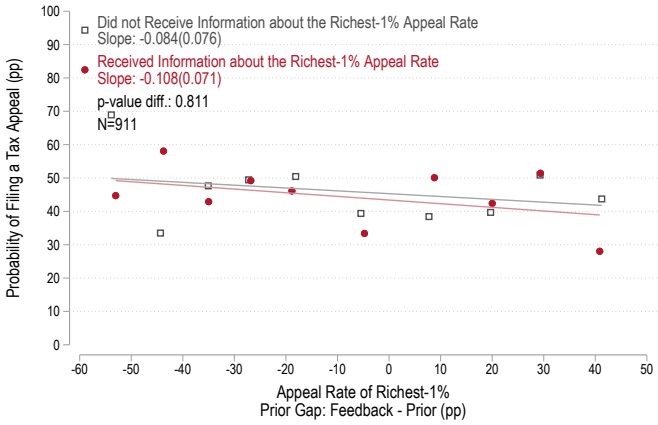
B. Expected Savings: Intention-to-Treat



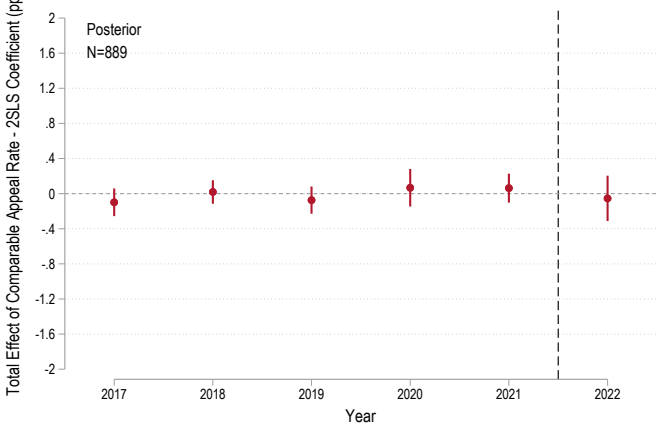
C. Richest-1% Appeal Rate: Event-Study



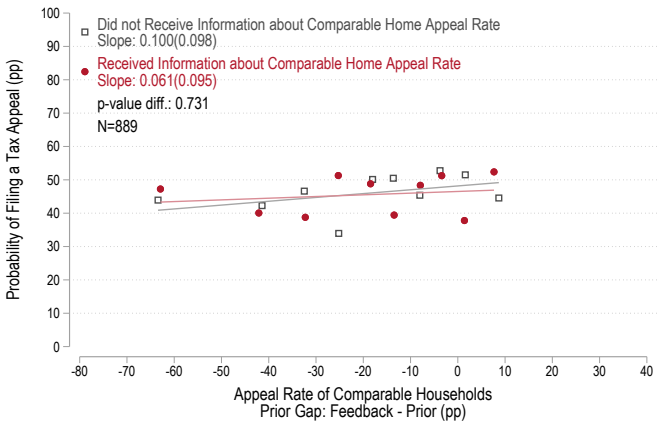
D. Richest-1% Appeal Rate: Intention-to-Treat



E. Comparable Appeal Rate: Event-Study



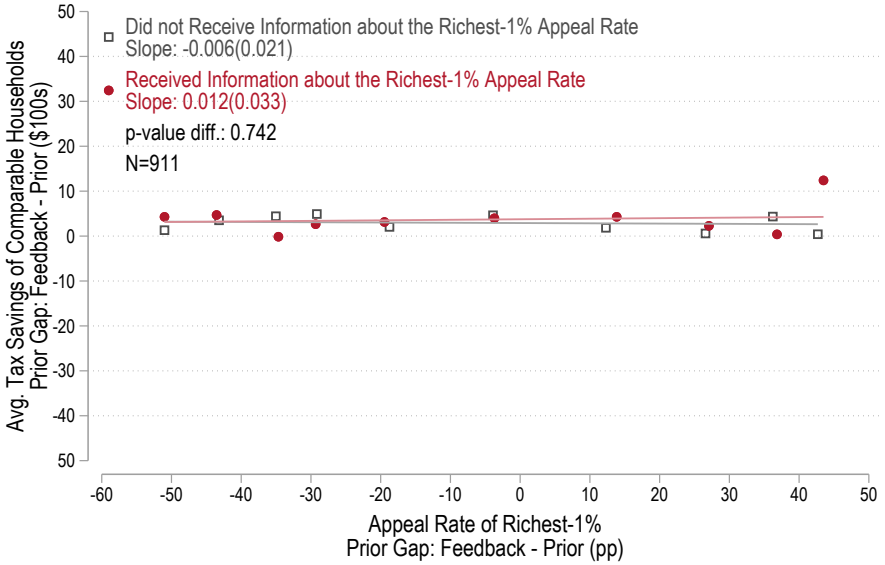
F. Comparable Appeal Rate: Intention-to-Treat



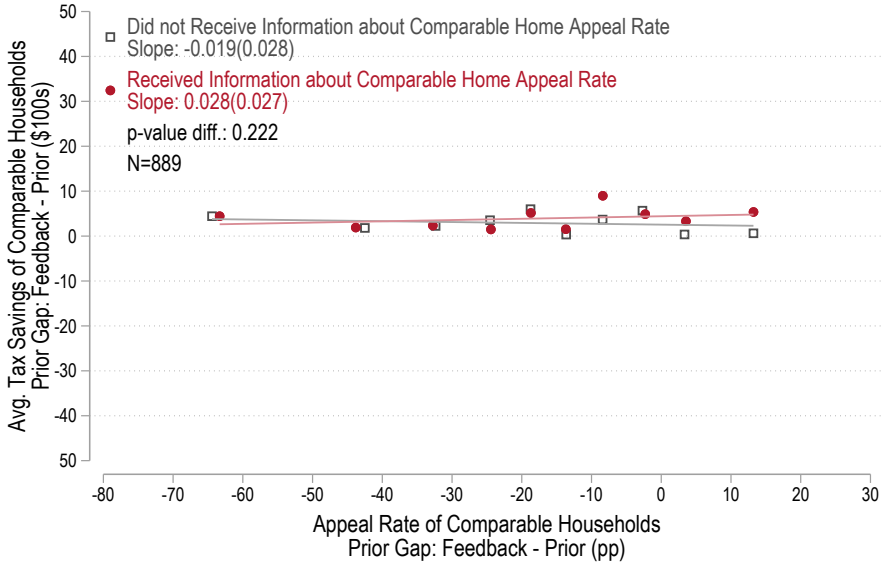
Notes: The left panels show event-study results on how posterior beliefs affect households' appeal probability (y-axis). The circles represent 2SLS estimates from Equation (5), with 90% confidence intervals based on robust standard errors. The vertical dashed line separates the post-treatment year (2022) from the pre-treatment years (2017-2021). The right panels show binned scatterplots of the intention-to-treat effects. Red circles represent treatment and gray squares control. The x-axes correspond to the difference between the observed value and the prior belief. Each line corresponds to an OLS binned scatterplot regression. Panel B controls for priors about the richest-1% and comparable homes, the comparable appeal rate information gap and whether the homestead cap is binding. Panel D controls for comparable home appeal rate priors and information gaps, and expected savings priors and information gaps interacted with whether the homestead cap is binding. Panel F controls for richest-1% appeal rate priors and information gaps, and expected savings priors and information gaps interacted with whether the homestead cap is binding.

Figure 8: Effects of Information about the Richest-1% and Comparable Appeal Rates on Expected Savings

A. Effect of Richest-1% Appeal Rate on Expected Savings



B. Effect of Comparable Appeal Rate on Expected Savings



Notes: Panel A shows how respondents update their beliefs about expected savings using a binned scatterplot. The x-axis reports the difference between the actual richest-1% appeal rate and respondents' prior beliefs about the richest-1% appeal rate. The y-axis reports the difference between posterior and prior beliefs about the expected savings from appealing (i.e., belief updating). Red circles (gray squares) represent the average update within each bin for the group of homeowners that were selected (were not selected) into the receiving information about the richest-1% appeal rate. Each line corresponds to the fitted values from a linear regression. The coefficients associated with the gap variable are reported in the upper left corner, as well as their robust standard errors (in parentheses), the p-value of the difference in the slopes, and the number of observations included in the analysis. Panel B shows how respondents update their beliefs about expected savings using a binned scatterplot. The x-axis reports the difference between the actual comparable appeal rate and respondents' prior beliefs about the comparable appeal rate. The y-axis reports the difference between posterior and prior beliefs about the expected savings from appealing (i.e., belief updating). Red circles (gray squares) represent the average update within each bin for the group of homeowners that were selected (were not selected) into the receiving information about the comparable appeal rate. Each line corresponds to the fitted values from a linear regression. The coefficients associated with the gap variable are reported in the upper left corner, as well as their robust standard errors (in parentheses), the p-value of the difference in the slopes, and the number of observations included in the analysis.

Table 1: Descriptive Statistics and Randomization Balance

		Treatment: Received information about Richest-1% | Comparable | Savings								
	All	N | N | N	Y | N | N	N | Y | N	N | N | Y	Y | Y | N	Y | N | Y	N | Y | Y	Y | Y | Y	p-value
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A. Administrative Data:										
2022 Home Value ($1,000)	448.006	447.368	455.159	454.253	450.997	445.348	441.955	431.548	456.057	0.857
	(4.224)	(12.704)	(11.768)	(11.400)	(12.605)	(11.073)	(11.311)	(11.780)	(13.001)	
2022 Property Tax Amount ($1,000s)	9.697	9.582	9.906	9.856	9.758	9.691	9.551	9.316	9.871	0.718
	(0.088)	(0.252)	(0.253)	(0.240)	(0.251)	(0.241)	(0.235)	(0.238)	(0.267)	
Homestead Exemption (%)	87.014	87.321	86.207	85.563	84.855	89.540	87.342	88.940	86.422	0.393
	(0.554)	(1.629)	(1.555)	(1.621)	(1.694)	(1.401)	(1.529)	(1.507)	(1.592)	
Homestead Exemption and Binding Cap (%)	66.205	69.378	64.097	65.180	62.138	69.247	66.034	67.972	65.948	0.253
	(0.780)	(2.257)	(2.163)	(2.197)	(2.292)	(2.113)	(2.178)	(2.242)	(2.202)	
2021 Owner Appeal Rate (%)	9.318	10.526	9.128	8.068	8.463	8.159	10.759	8.525	10.991	0.565
	(0.479)	(1.503)	(1.298)	(1.256)	(1.315)	(1.253)	(1.425)	(1.342)	(1.454)	
2020 Owner Appeal Rate (%)	17.169	16.746	16.836	16.561	15.145	16.946	19.409	16.590	18.966	0.739
	(0.622)	(1.828)	(1.687)	(1.715)	(1.694)	(1.718)	(1.819)	(1.788)	(1.822)	
2019 Owner Appeal Rate (%)	11.464	10.048	10.953	9.130	13.363	12.134	11.392	12.212	12.500	0.538
	(0.525)	(1.472)	(1.408)	(1.329)	(1.608)	(1.495)	(1.461)	(1.573)	(1.537)	
2018 Owner Appeal Rate (%)	10.052	8.373	11.359	8.280	9.800	10.251	9.916	12.673	9.698	0.399
	(0.496)	(1.356)	(1.431)	(1.271)	(1.405)	(1.389)	(1.374)	(1.599)	(1.375)	
Multiple Owners (%)	30.128	27.751	34.280	28.662	29.176	30.544	29.958	29.954	30.172	0.554
	(0.756)	(2.193)	(2.140)	(2.086)	(2.148)	(2.109)	(2.106)	(2.201)	(2.133)	
Living Area (1,000s Sq. Feet)	2.302	2.239	2.336	2.314	2.346	2.324	2.282	2.269	2.298	0.625
	(0.014)	(0.041)	(0.039)	(0.040)	(0.043)	(0.041)	(0.039)	(0.042)	(0.041)	
Number of Bedrooms	3.419	3.366	3.448	3.363	3.463	3.458	3.414	3.472	3.362	0.068
	(0.012)	(0.037)	(0.033)	(0.035)	(0.036)	(0.034)	(0.035)	(0.034)	(0.036)	
Number of Baths	2.287	2.242	2.284	2.272	2.303	2.331	2.291	2.279	2.291	0.816
	(0.012)	(0.036)	(0.034)	(0.035)	(0.035)	(0.035)	(0.033)	(0.034)	(0.035)	
Panel B. Survey Data:										
Republican (%)	32.723	31.969	34.534	30.425	31.765	34.792	33.187	32.673	32.135	0.880
	(0.794)	(2.361)	(2.191)	(2.179)	(2.261)	(2.231)	(2.210)	(2.336)	(2.216)	
Democrat (%)	49.170	46.803	50.212	48.546	48.941	51.204	47.473	51.238	48.764	0.870
	(0.846)	(2.527)	(2.304)	(2.367)	(2.428)	(2.341)	(2.344)	(2.490)	(2.372)	
Female (%)	41.726	43.401	42.947	40.133	40.047	41.215	44.565	41.872	39.644	0.770
	(0.831)	(2.500)	(2.274)	(2.311)	(2.374)	(2.295)	(2.320)	(2.451)	(2.311)	
Age	47.005	46.916	46.989	47.257	47.173	47.406	46.515	46.456	47.269	0.879
	(0.188)	(0.575)	(0.538)	(0.516)	(0.526)	(0.525)	(0.523)	(0.545)	(0.516)	
Race: White (%)	55.436	54.569	56.632	54.102	51.991	58.568	55.435	58.867	53.229	0.387
	(0.838)	(2.512)	(2.276)	(2.349)	(2.421)	(2.297)	(2.320)	(2.445)	(2.357)	
Education: Grad. Degree (%)	84.246	86.548	85.263	84.701	85.948	82.863	80.652	84.975	83.519	0.297
	(0.614)	(1.721)	(1.628)	(1.697)	(1.684)	(1.757)	(1.844)	(1.776)	(1.753)	
Observations	3,681	418	493	471	449	478	474	434	464	

Notes: Standard errors are reported in parentheses. The statistics in Panel A are based on administrative records available on the DCAD website. The statistics in Panel B are based on survey responses. Column (1) is based on the entire subject pool. Column (2) is based on homeowners not selected to receive any information (N | N | N). Column (3) is based on homeowners selected to receive information on the appeal rate of the richest-1% of homeowners in 2021. Column (4) is based on homeowners selected to receive information on the appeal rate of comparable homeowners in 2021. Column (5) is based on the expected savings of comparable homeowners in 2021. Column (6) is selected to receive information on the richest-1% appeal rate and comparable households appeal rate. Column (7) is selected to receive information on the richest-1% appeal rate and the expected savings of comparable homeowners in 2021. Column (8) is selected to receive the appeal rate of comparable homeowners and the expected savings of comparable homeowners in 2021. Column (9) receives all three pieces of information for homeowners in 2021. Column (10) reports the p-value of a test of equal means across the eight treatment groups.

Table 2: Causal Effects of Beliefs (2SLS Estimates)

	Owner Appeal (1)	Appeal Intention (2)	Agent Appeal (3)	Successful Appeal (4)	Appeal Ineq. is Fair (5)	Tax Appeals Always Just (6)
Panel A: Total Effects						
Expected Savings ($100s)	1.302**	-0.472	-0.105	1.156**	0.020*	0.554
	(0.570)	(0.612)	(0.253)	(0.551)	(0.012)	(0.423)
Richest-1% Appeal Rate	-0.034	-0.137	-0.020	-0.126	-0.005	0.111
	(0.142)	(0.136)	(0.077)	(0.138)	(0.003)	(0.107)
Comparable Appeal Rate	-0.054	0.101	-0.052	0.016	-0.002	0.262**
	(0.157)	(0.146)	(0.077)	(0.151)	(0.004)	(0.114)
Kleibergen-Paap F-Statistic:						
Expected Savings	12.868	12.785	12.868	12.868	14.794	12.785
Richest-1% Appeal Rate	561.011	561.011	561.011	561.011	528.353	561.011
Comparable Appeal Rate	179.349	179.349	179.349	179.349	161.714	179.343
Observations:						
Expected Savings	867	865	867	867	816	865
Richest-1% Appeal Rate	911	911	911	911	863	911
Comparable Appeal Rate	889	889	889	889	839	888
Panel B: Partial Effects						
Expected Savings ($100s)	0.583***	-0.062	0.047	0.496**	0.003	0.269
	(0.206)	(0.215)	(0.106)	(0.195)	(0.005)	(0.169)
Richest-1% Appeal Rate	0.126	0.123	0.000	0.075	-0.005***	0.090
	(0.086)	(0.080)	(0.048)	(0.083)	(0.002)	(0.066)
Comparable Appeal Rate	-0.072	0.051	-0.001	-0.071	0.006***	0.053
	(0.104)	(0.095)	(0.056)	(0.100)	(0.002)	(0.080)
Kleibergen-Paap F-statistic	15.106	15.027	15.106	15.106	15.196	14.986
Mean Outcome (Control)	45.933	62.919	7.416	35.885	0.000	83.014
Observations	3,681	3,675	3,681	3,681	3,502	3,673

Notes: Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Robust standard errors in parentheses. Panel A: Each column reports 2SLS estimates from three regressions using Equation 5 for expected savings, richest-1% appeal rate, and comparable appeal rate. Panel B: Each column reports 2SLS estimates from one regression using Equation 6. The outcome variable is indicated in the column title. Columns (1), (3), and (4) use outcome data from administrative tax records. Columns (2), (5), and (6) use survey responses. Column (5) is standardized to have a mean of zero and a standard deviation 1 in the control group. Column (6) is an indicator equal to 100 if the subject said that tax appeals were always completely justified.

Online Appendix (For Online Publication Only)

"Does Tax Avoidance Trickle Down? Evidence from a Field Experiment"

Holz, Perez-Truglia and Zentner

February 6, 2026

A Further Details and Results

A.1 Details about the Field Survey Sample

Table A.1 reports summary statistics for the 100,615 households invited to participate in the study and compares the characteristics of those in the subject pool with those who are not. The latter group includes households who did not respond to the survey as well as a minority who began the survey but were excluded for various reasons, such as not completing the main module, reporting outlier beliefs, or having appealed before answering the survey. There are some differences between the two groups, although typically modest in magnitude. For instance, households in the subject pool have an average property value of $448,000, compared with $374,000 for those outside the pool. The largest difference, as discussed in Section 3.5, is in the appeal rate: 9.3% of households in the subject pool filed a tax appeal in the year prior to the experiment (2021), compared with 4.6% among those not in the subject pool.

A.2 Feedback Provided in the Information-Provision Experiment

Table A.2 summarizes the feedback provided in the information-provision experiment. The first column lists the home value bins, and the second column reports the share of subjects in each bin. The third column shows that all subjects received the same feedback (49%) regarding the appeal rate of the richest 1%. The fourth column reports the appeal rate of comparable homes. Because higher-value homes are more likely to appeal, the feedback naturally depends on the subject's home value bin. A majority of subjects (53%) own homes valued between $150,000 and $350,000, for which the feedback is relatively similar, between 6% and 13%. For subjects at the top of the distribution, however, the feedback is much higher—for instance, for those in the 90th percentile, with homes worth $650,000–$699,999, the feedback was 30%. The fifth column reports the feedback on the average tax savings of comparable homes. Since the property tax amount increases with home value, it is natural for the expected dollar savings to vary across bins of home value. For the majority of subjects (53%), in homes between $150,000 and $350,000, the feedback ranges from $294 to $400. But

for wealthier households, the dollars savings can be larger—for example, for subjects in the 90th percentile, with homes worth $650,000–$699,999, the feedback on tax savings was $750.

A.3 Decision Trees used for the Willingness to Pay Elicitation

Section 3.3 explains the dynamic choices subjects faced in the question on willingness to pay to help someone file a tax appeal or to prevent that help. Figure A.1 provides the complete decision trees used in this question. Subjects first decided whether they wanted to help other taxpayers appeal their taxes. Those who chose to help followed the decision tree in Panel A, while those who chose not to help followed the decision tree in Panel B.

A.4 Details and Additional Results about the Attitudes Survey

We recruited 600 U.S. homeowners for the Attitudes Survey. Table A.3 compares the average characteristics of respondents in the Attitudes Survey and the field experiment. Because the two samples were recruited through different methods, their characteristics are not expected to align closely. Indeed, the differences in average characteristics are statistically significant across the board, though often modest in magnitude. Overall, respondents in the field experiment tend to be older, more likely to be male, and less left-leaning.

In addition to the questions on the acceptability of tax avoidance and evasion behaviors discussed in the main text, the survey included questions on awareness of and prior engagement with these behaviors. Specifically, we asked respondents two questions: "Before taking this survey, were you aware of any of these tax reduction strategies?" and "Have you ever used any of these strategies to lower your taxes?" We listed all the tax avoidance and evasion behaviors from the survey, allowing respondents to select all that applied.

Panel A of Figure A.2 presents the results on awareness. At one end, the tax behaviors with the highest awareness are the two forms of tax evasion—"Paying Under the Table" and "Under Report Income"—with a majority (about 60%) of respondents reporting awareness. At the other end, more obscure strategies used primarily by the ultra-wealthy—"Buy, Borrow, Die" and "Carried Interest"—were known by only a minority (about 20%). "Property Tax Appeals" falls toward the higher end of the awareness spectrum, with 44.5% of respondents reporting awareness of this strategy.

Panel B of Figure A.2 shows the responses on prior engagement. These results should be interpreted with caution—due to social desirability bias, respondents likely under-reported such behavior. "Property Tax Appeals" ranks near the top, with 24.1% of respondents reporting that they had filed a property tax appeal in the past.

Lastly, at the end of the survey we included an optional question designed to flag respondents who may have used an AI agent to complete it. To the best of our knowledge, at the time of data collection no standard or effective method existed for detecting AI agents in online surveys, so we devised our own approach. Specifically, we asked respondents to reproduce copyrighted material (the opening credits to *Star Wars*), noting that AI models are typically trained to refuse such requests. Approximately 78% of respondents answered this question, suggesting they were unlikely to have used an AI agent. For the remaining respondents, AI involvement is uncertain—some genuine participants may simply have chosen to skip the question. In any case, the results remain essentially unchanged when we exclude this 22% of the sample.

A.5 Outliers Beliefs

In total, 248 observations are excluded from the baseline specification because at least one of their prior beliefs was classified as an outlier. We now provide more detail on this classification. Figure A.3 shows the distribution of prior misperceptions, with each panel corresponding to one of the three beliefs. Outliers are shown in red and non-outliers in gray.

Panel A of Figure A.3 displays the distribution of prior misperceptions for the richest-1% appeal rate. The distribution is fairly uniform, with no obvious outliers. This is intuitive: since the true value is close to 50% and responses are bounded between 0% and 100%, there is limited scope for extreme misperceptions. In turn, Panel B of Figure A.3 shows the distribution for the comparable appeal rate. Here the distribution takes on a bell shape, with a small number of responses at the far left and far right, possibly reflecting outliers due to typos or misunderstandings. We aimed to classify as outliers the misperceptions in the bottom 1% and top 1%, but could not achieve this exactly due to ties. We classified as outliers those in the bottom 0.8% and top 0.9%, which corresponds to 139 subjects.

Panel C of Figure A.3 shows the distribution of misperceptions about expected tax savings. For ease of interpretation, this figure reports average expected savings as a share of the tax amount of the corresponding group of comparable homes. Because savings cannot be negative, underestimation is naturally bounded. Overestimation, by contrast, has no upper limit, leaving considerable scope for extreme responses. Indeed, Panel C shows that some individuals vastly overestimated average tax savings—for example, some respondents reported values as high as 90%. These extreme responses are most likely reflecting typos or misunderstandings. Given the natural boundary on underestimation, we focus on the bottom 2% rather than trimming from both tails. While we aimed for the bottom 2%, exact targeting was not possible due to ties. We classified as outliers the bottom 1.87% of the distribution, corresponding to 116 individuals.

A.6 Additional Results on Belief Updating

As discussed in Section 2, households without a binding cap are expected to update their beliefs about expected savings more than those with a binding cap. Figure A.4 illustrates this comparison: Panel A shows updating for households without a binding homestead cap, while Panel B shows updating for households with a binding homestead cap. The difference in the slopes when the feedback is shown vs. not shown in Panel A is 0.67, much larger than the difference in slopes of 0.25 in Panel B. These results confirm our expectations, showing substantially greater belief updating when the cap is not binding than when it is binding.

The information provision about one belief may influence that belief directly but it may also indirectly influence other beliefs, which is known as cross-learning. Figure 8 in the main text shows that learning about the appeal rates of the richest 1% or of comparable households does not lead households to revise their expectations about their own savings. We now test other potential forms of cross-learning in the information provision experiment. In Panel A of Figure A.6, the x-axis displays the difference between the prior belief and the feedback in expected savings, and the y-axis displays the belief update for the comparable appeal rate. The results in Panel A indicate that the expected-savings treatment produces no meaningful shifts in the comparable appeal rate beliefs. Similarly, Panel B shows that the expected-savings treatment produces no meaningful shifts in the richest-1% appeal rate beliefs. By contrast, Panels C and D reveal statistically significant, though modest, cross-learning between the two appeal-rate beliefs: updating beliefs about the richest-1% leads to small adjustments in beliefs about comparable homes, and vice versa. For example, in Panel D of Figure A.6, the difference in slopes between treatment and control groups is small $0.114\ (= 0.185 - 0.071)$ but statistically significant (p-value < 0.001). This evidence suggests that learning the richest-1% appeal rate is 1 pp higher than initially believed increases the perceived appeal rate of comparable homes by about 0.11 pp.

A.7 Additional Descriptive Statistics

In the field experiment, we included two questions that serve as dependent variables: one measuring the perceived fairness of the wealth gap in appeal rates, and another measuring how justifiable it is to file a tax appeal. Figure A.7 shows the distribution of responses to each question, with one panel per outcome.

A.8 Further Results and Robustness Checks

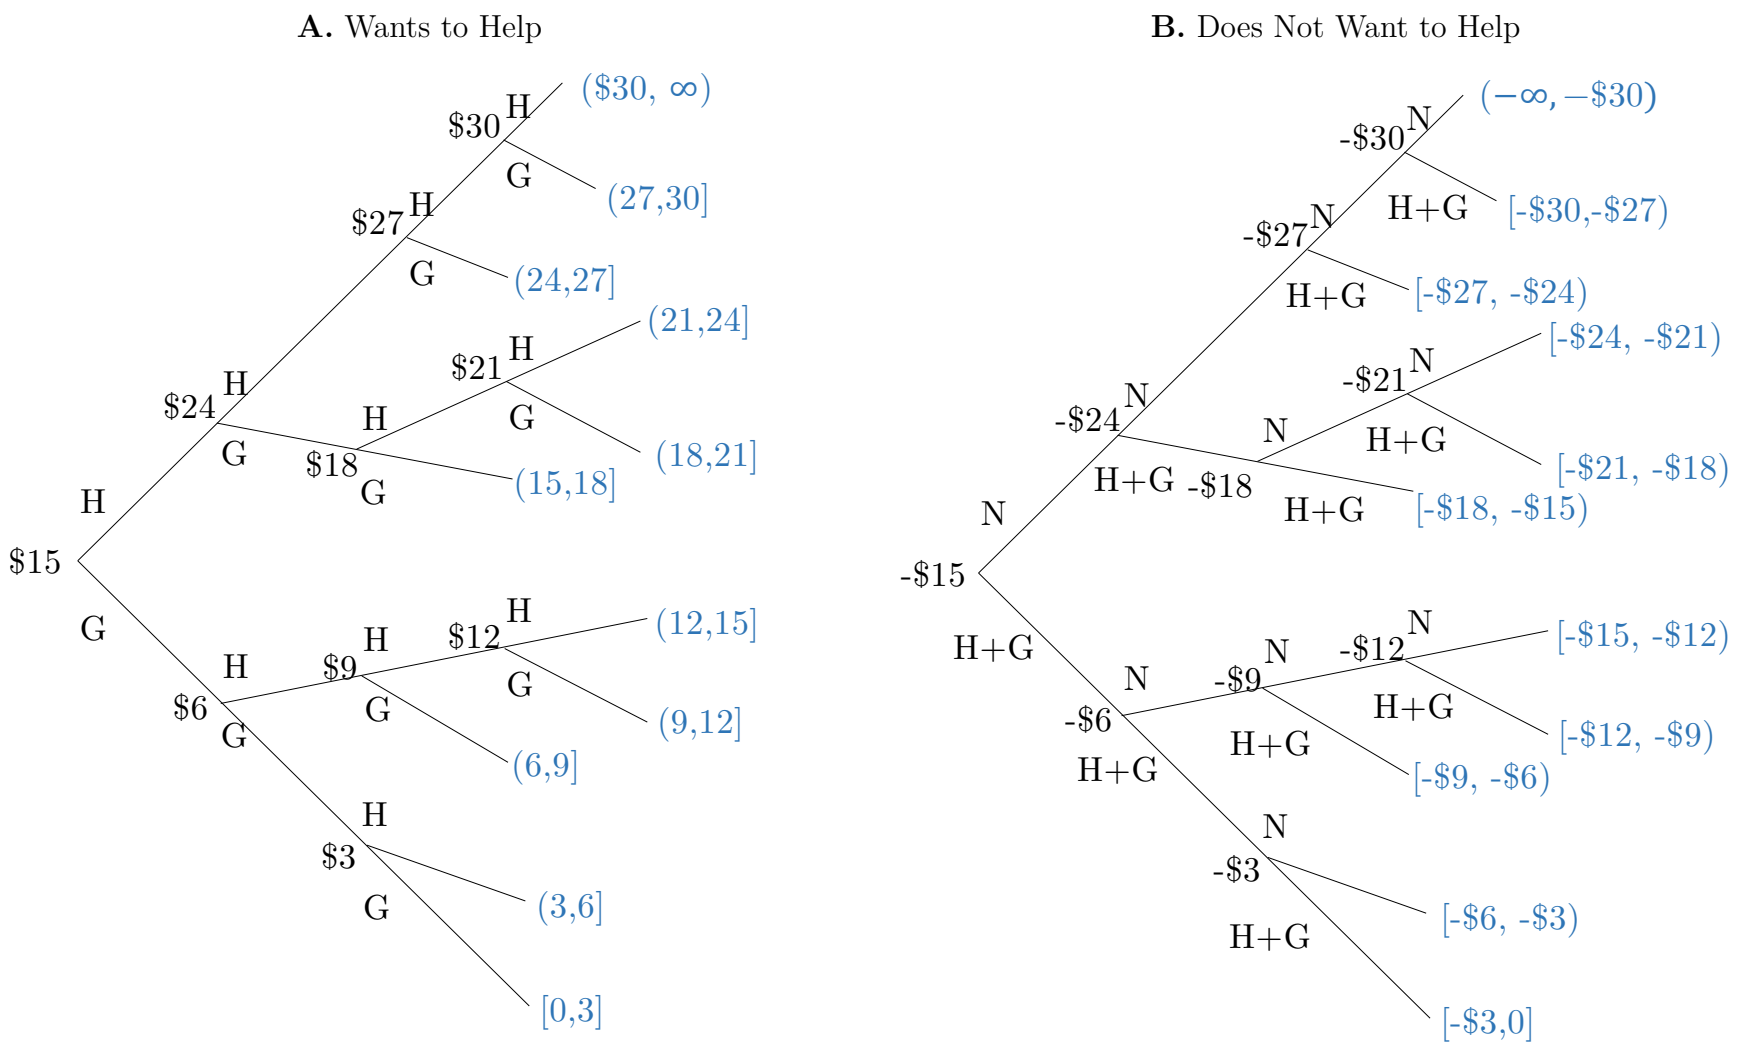
Figure 7 in the main text presents the event-study analysis and binned scatterplot analysis for the total effects. Figure A.8 replicates the same analysis, but corresponding to the partial

effects instead of the total effects. For example, for expected savings (Panel A of Figure A.8), the coefficient estimate is statistically significant in 2021, the year prior to our intervention. The negative sign of this estimate suggests that our 2022 estimates may underestimate the true effect. Second, we find a statistically significant coefficient estimate in 2020 for the appeals rate of the richest-1%, which is of the same sign and similar magnitude as the 2022 estimate. Finally, we find a statistically significant coefficient estimate in 2017 for appeals of comparable homes.

Regarding the event-study graphs, there are some subtle differences. The estimated "effects" in pre-treatment years are, as expected, always close to zero. The key difference, however, is that some of these pre-treatment effects, although small, are statistically significant—specifically, the coefficient for expected savings in 2021, the coefficients for the richest-1% appeal rate in 2020 and 2018, and the coefficient for the comparable appeal rate in 2017. Because we conduct many falsification checks, it is not surprising that some estimates appear spuriously significant purely by chance. However, out of 15 pre-treatment coefficients we would expect about 1.5 to be statistically significant at the 10% level by chance, but we observe 4. Because of this minor anomaly in the partial effects, we have slightly greater confidence in the total effects specification than in the partial effects specification.

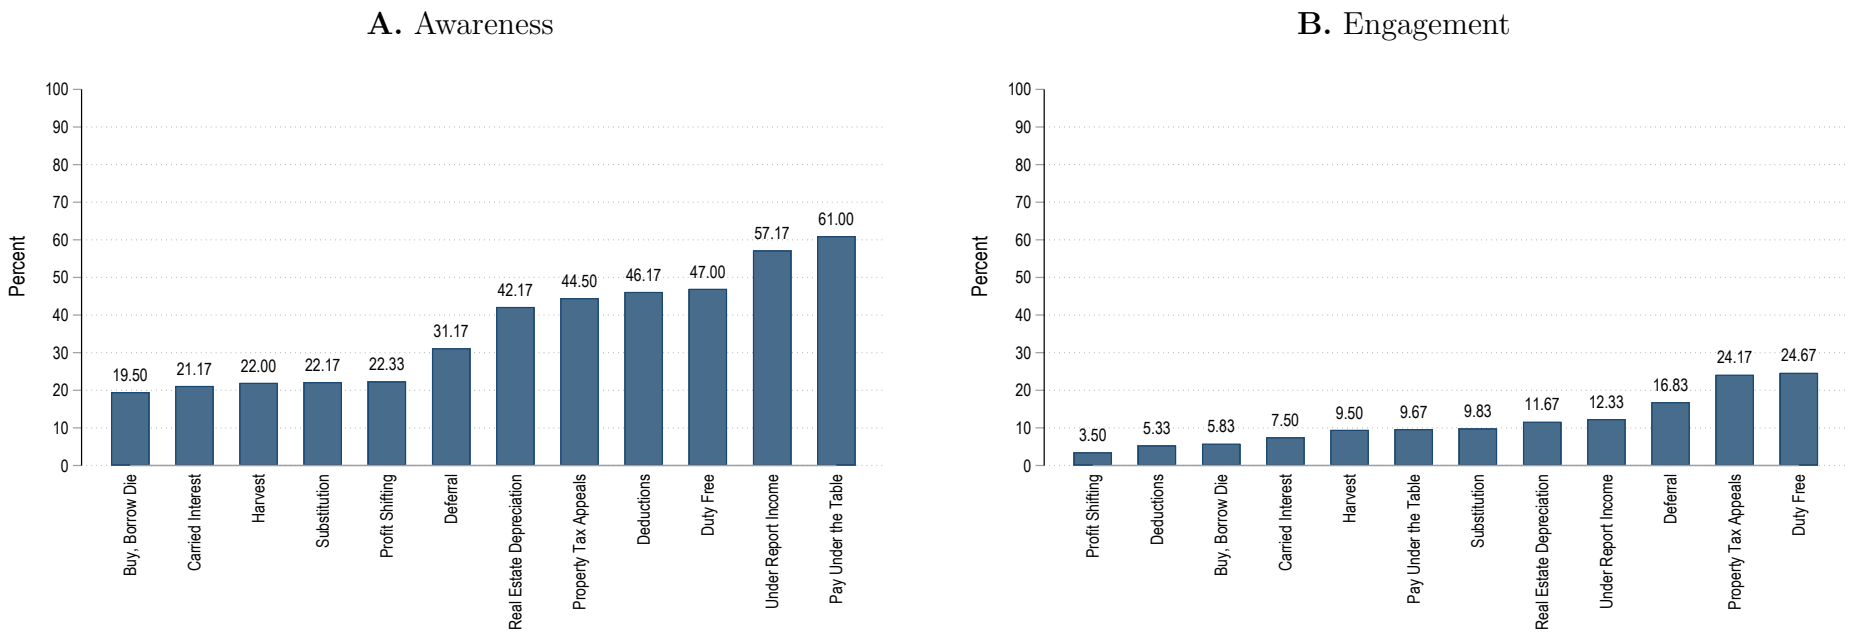
Table A.7 presents further robustness checks. Column (1) reproduces the baseline specification from Table 2, while the other columns report results from alternative specifications. Column (2) includes additional control variables—whether the subject previously appealed, and a set of subject characteristics: age, gender, college graduation and political party. Because these characteristics were measured at the end of the survey, and some respondents completed only the main module rather than the full survey, including these controls slightly reduces the sample size. Column (3) includes individuals with outlier misperceptions of expected savings, and column (4) includes those with outlier misperceptions of comparable appeal rates. Column (5) excludes subjects who failed the attention check. Across all specifications, the results remain consistent. The only notable difference is that in column (3), the total effect of tax savings is larger than in column (1), but due to much larger standard errors, it becomes statistically insignificant. This is expected, as including extreme outliers typically reduces precision. By contrast, relative to column (1), the total effects in column (4) are larger in magnitude and remain statistically significant.

Figure A.1: Decision Trees for Willingness to Pay Measurement



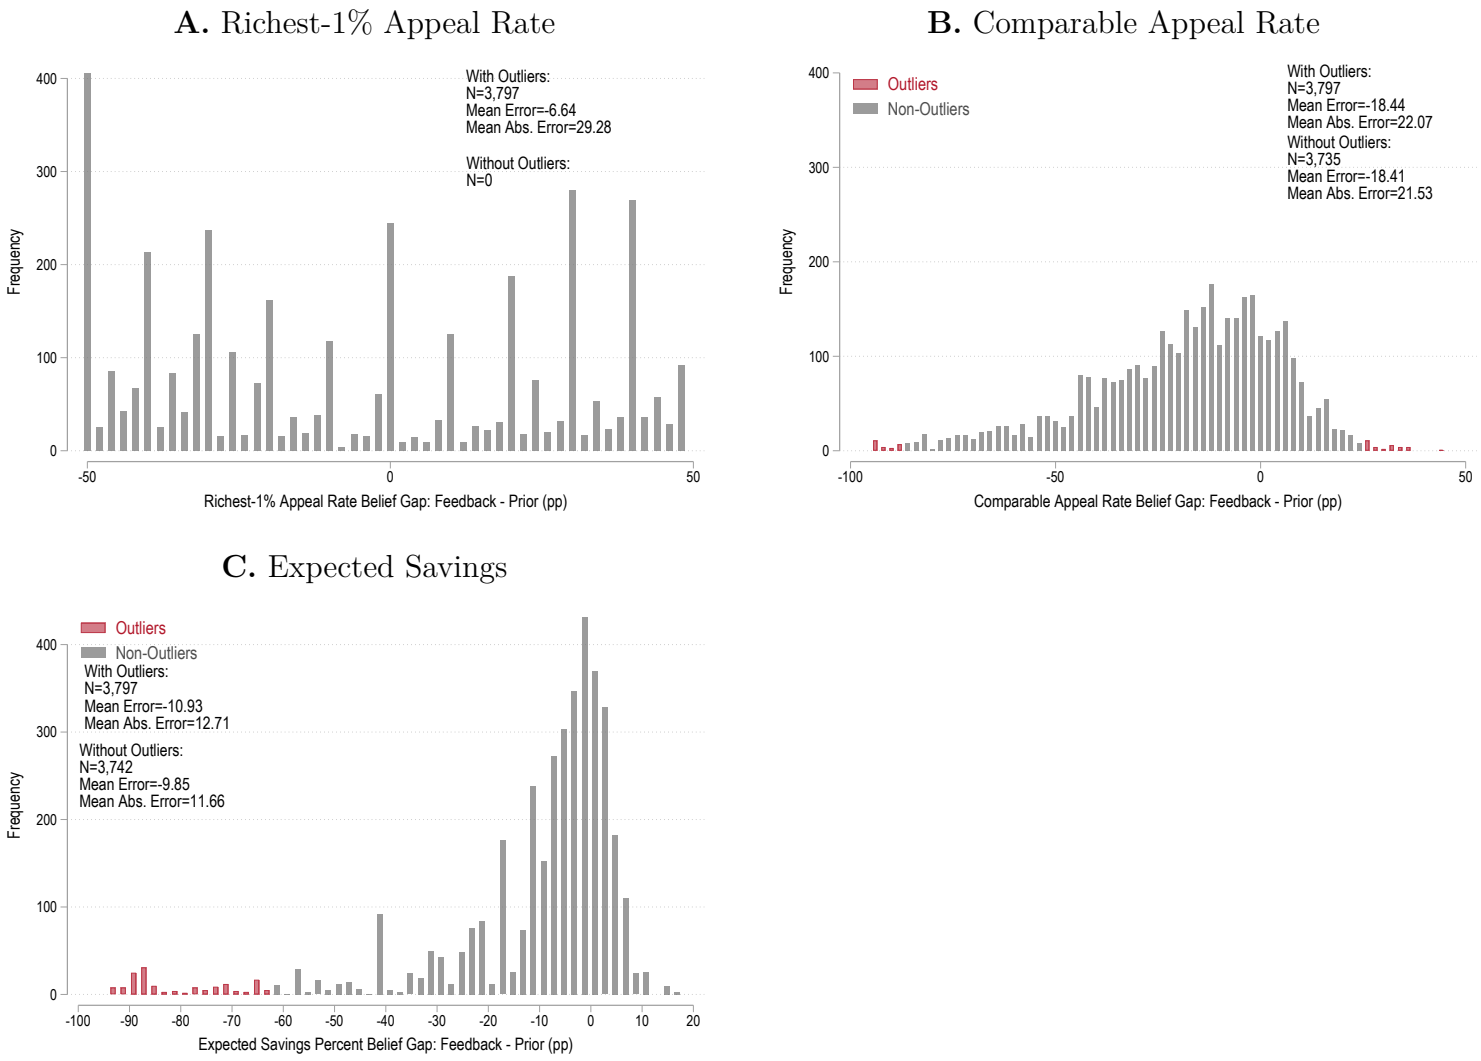
Notes: This figure displays the decision trees that subjects faced during the willingness to pay to help others appeal their taxes. In this figure, "G" represents the subject's decision to choose the gift card, "H" represents the subject's decision to have us help the taxpayer, and "N" represents the subject's decision to have us not help the taxpayer. All subjects began by deciding whether to help the other taxpayer. If they chose to help the other taxpayer, they went through the decision tree in Panel A. If they chose not to help the other taxpayer, they went through the decision tree in Panel B.

Figure A.2: Awareness and Engagement of Tax Avoidance/Evasion Strategies



Notes: This figure uses data from the online perceptions survey. Panel A displays the percentage of subjects who were previously aware of each strategy. Panel B displays the percent of subjects who had previously engaged in each strategy.

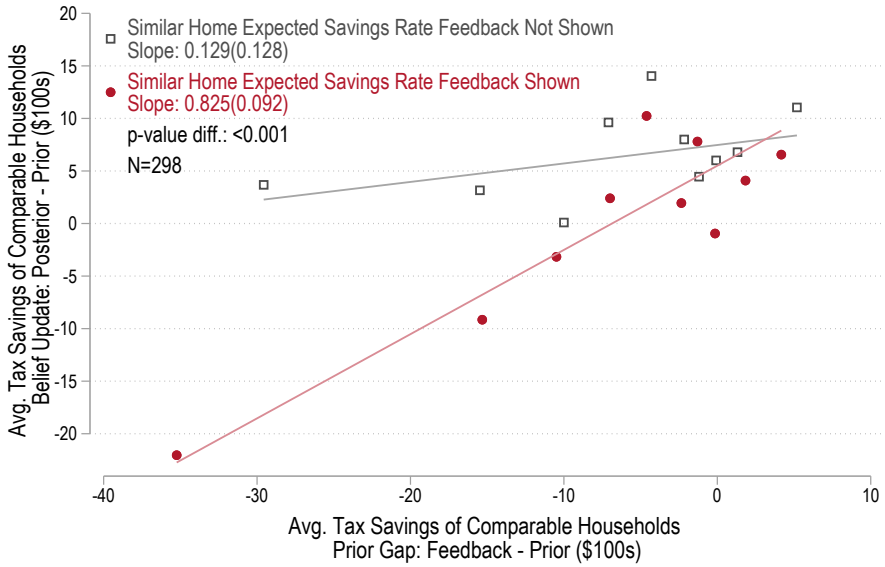
Figure A.3: Distribution of the Appeal Rate and Expected Savings Gaps



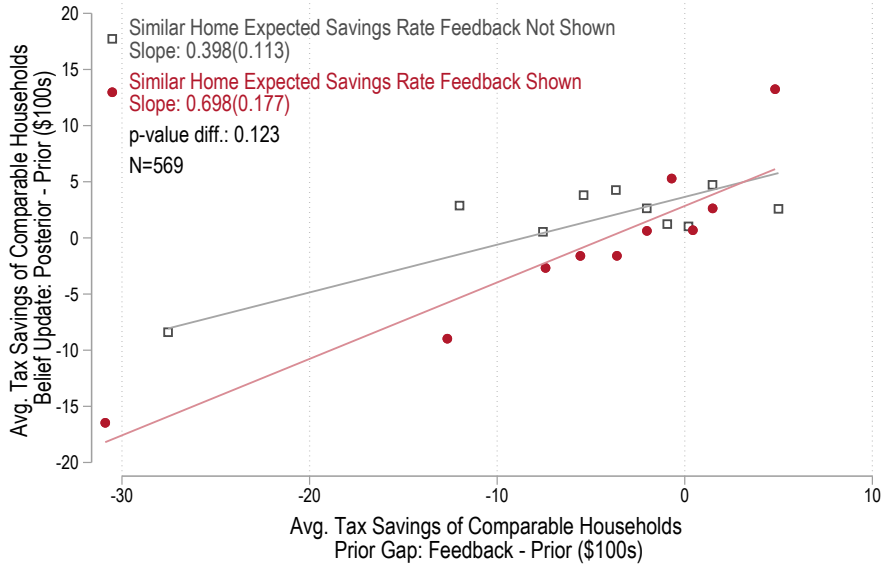
Notes: This figure displays the distribution of prior beliefs in gray along with the responses marked as outliers in red. Panel A displays the distribution of prior beliefs for the richest-1% appeal rates. We do not classify any observations as outliers in this figure. Panel B displays the distribution of prior beliefs for the comparable homes appeal rates. Panel C displays the distribution of prior beliefs about expected savings.

Figure A.4: Updating on Tax Savings With and Without a Binding Homestead Cap

A. Without Binding Homestead Cap

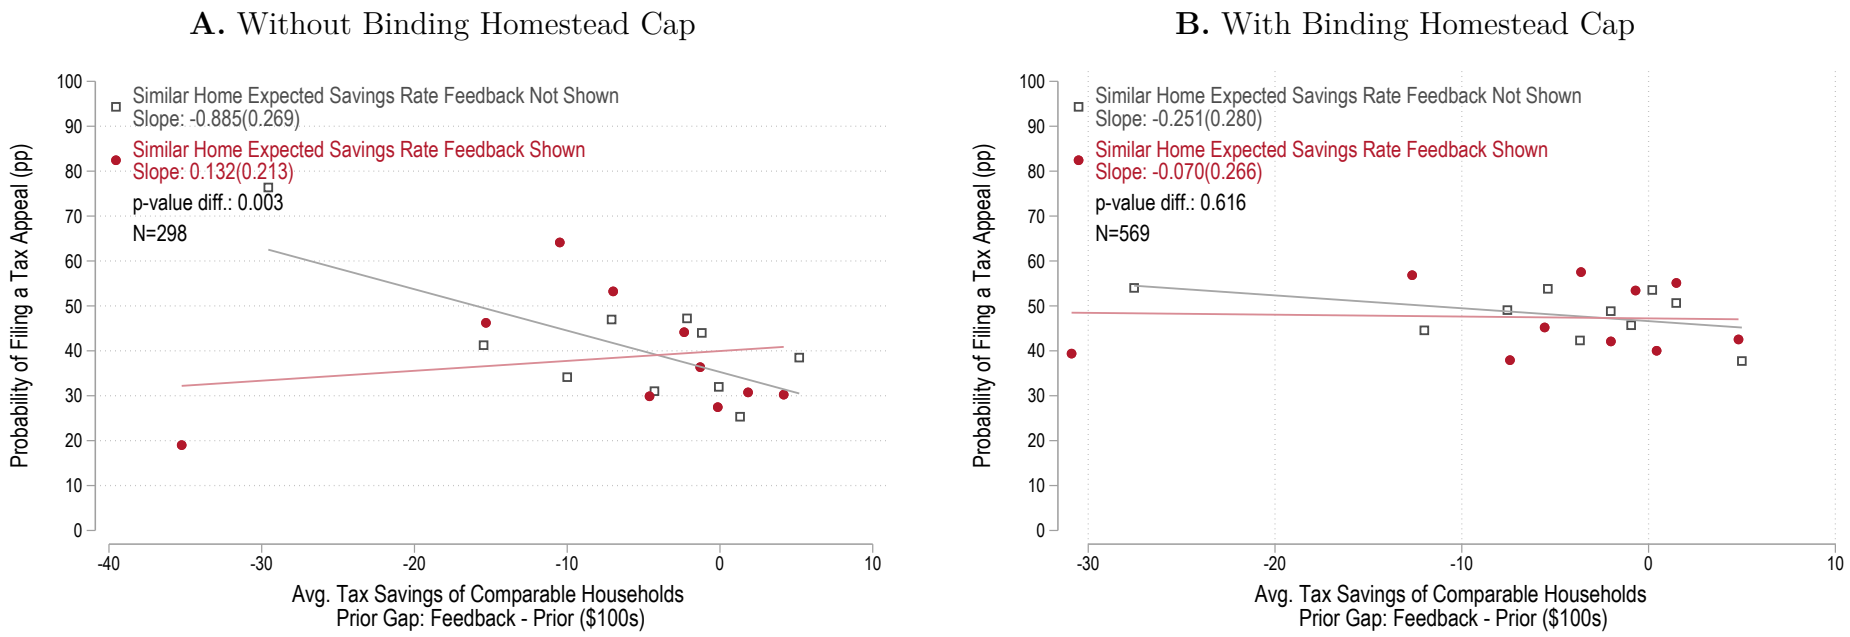


B. With Binding Homestead Cap



Notes: This figure is analogous to Panel B of Figure 6 in the main text, but separating the subjects by whether or not they have a binding homestead cap.

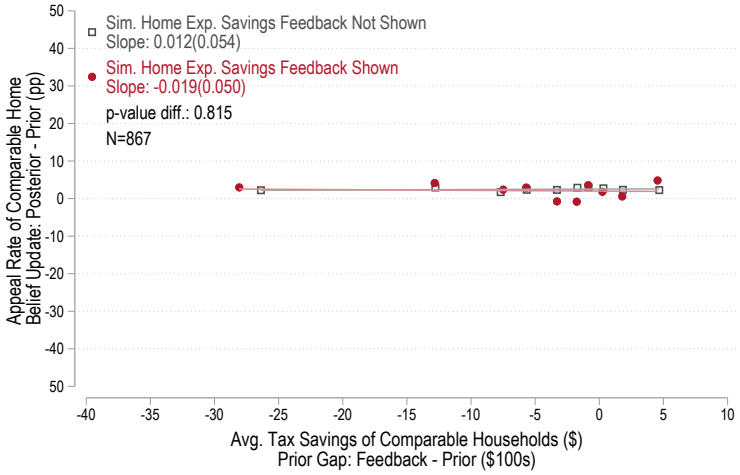
Figure A.5: Intention-to-Treat Effect of Savings with and withing a Binding Homestead Cap



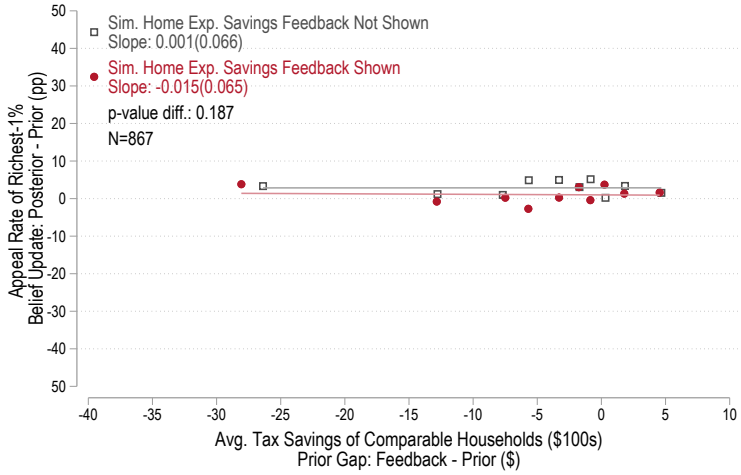
Notes: This figure is analogous to Panel B of Figure 7 in the main text, but separating the subjects by whether or not they have a binding homestead cap.

Figure A.6: Additional Tests of Cross-Learning in the Information-Provision Experiment

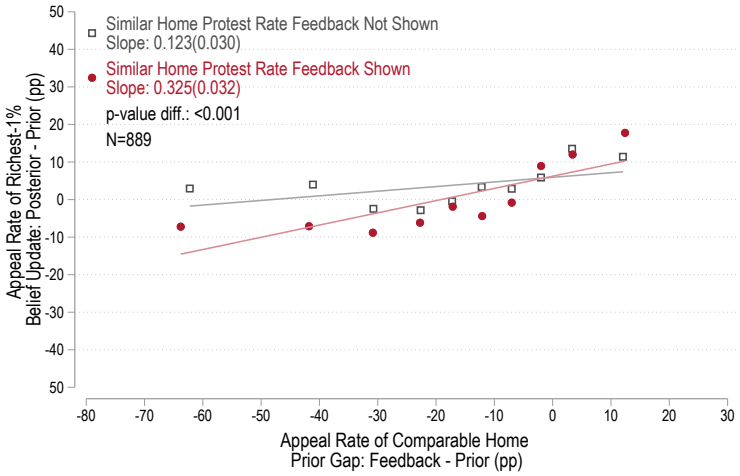
A. From Savings to Comparable Appeal Rate



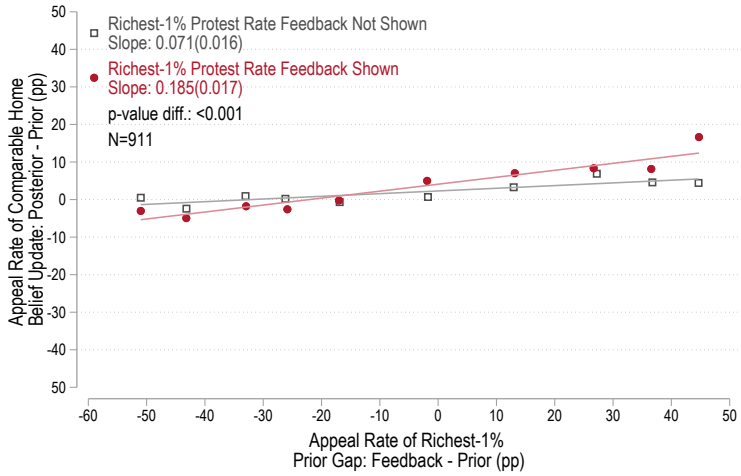
B. From Savings to Richest-1%



C. From Comparable Appeal Rate to Richest-1%

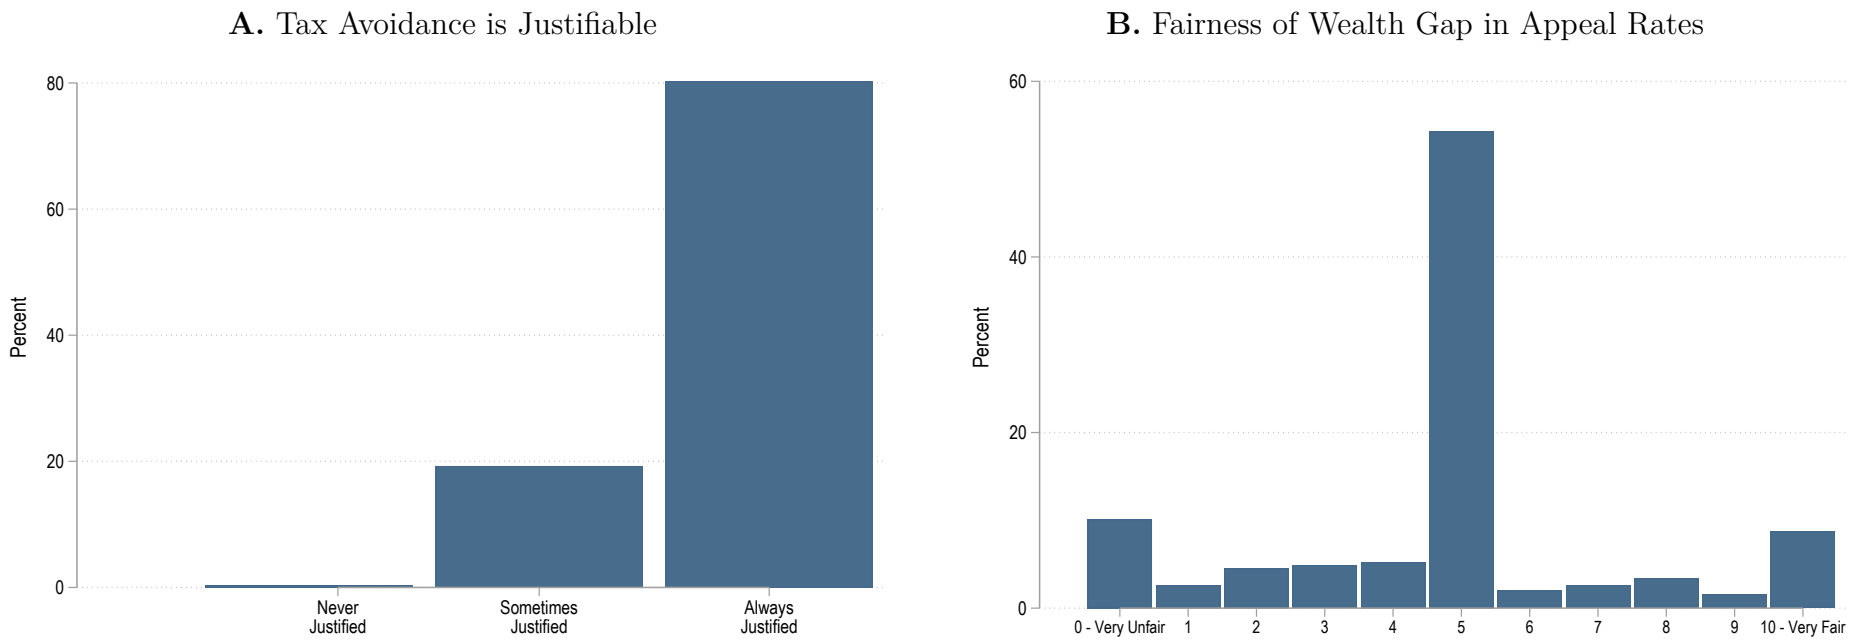


D. From Richest-1% to Comparable Appeal Rate



Notes: This figure is analogous to Figure 8 in the main text, but testing other cross-learning possibilities in the information provision experiment.

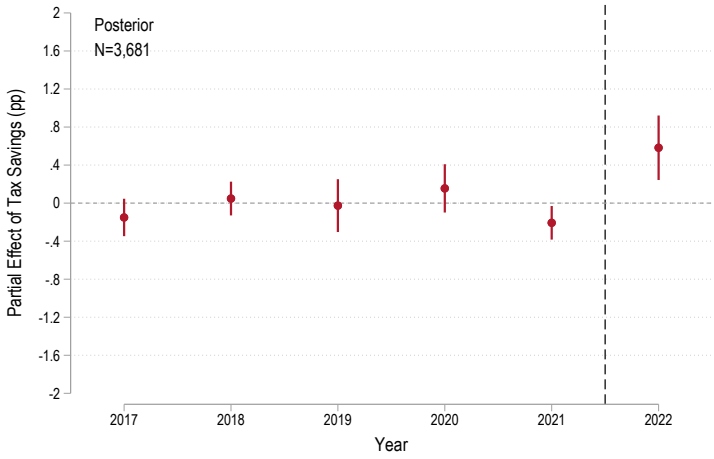
Figure A.7: Additional Descriptive Analysis



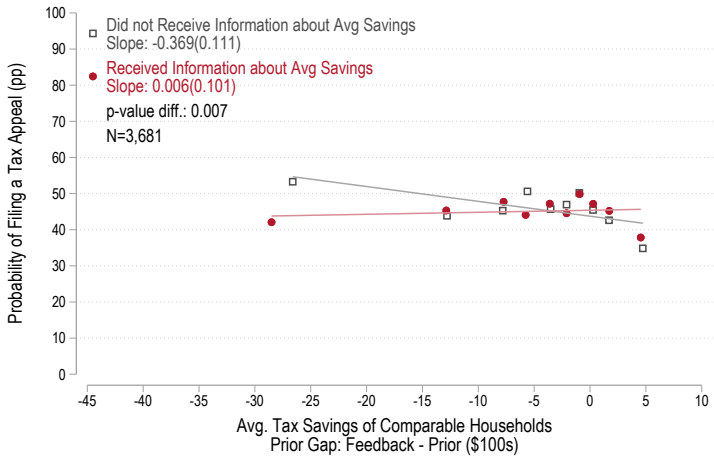
Notes: Panel A shows the distribution of responses to the question "How justifiable is it to lawfully reduce your tax bill if you have a chance?" Panel B shows the distribution of responses to "You said that [Richest-1% appeal rate posterior]% of the richest households file their tax protest, while [Comparable appeal rate posterior]% of households like yours file a protest. Do you think this is fair or unfair?" The posterior beliefs are subject-specific and those elicited earlier in the experiment. The x-axis in Panel A represents a justifiability scale, and in Panel B it represents a 0–10 fairness scale. The y-axis in both panels shows the fraction of subjects selecting each possible choice.

Figure A.8: 2SLS Partial Effects: Event-Study Analysis and Reduced-Form Binned Scatterplots

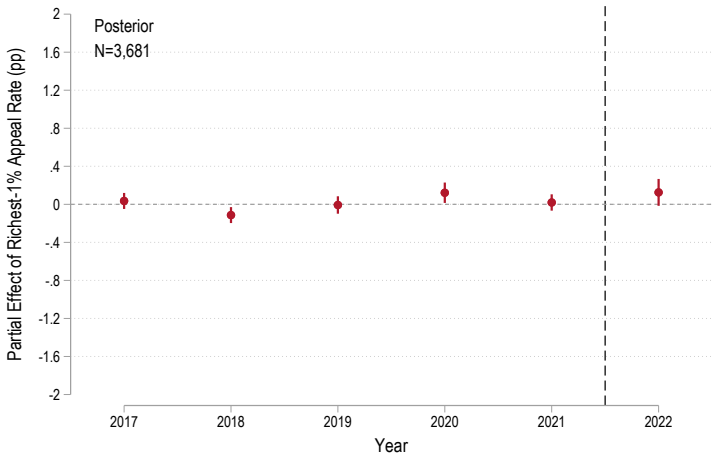
A. Expected Savings: Event-Study



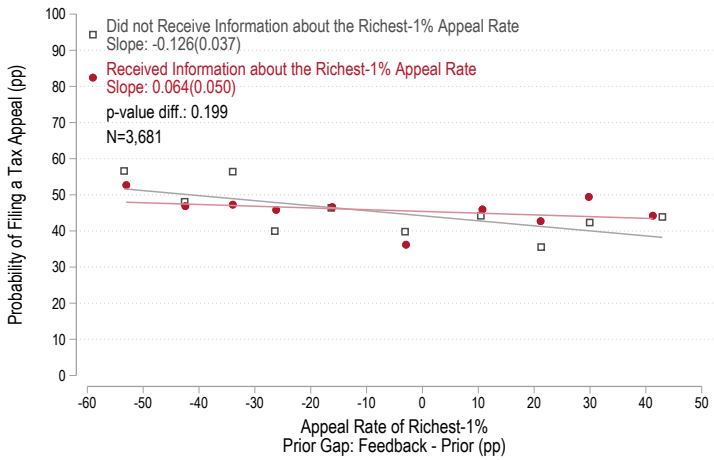
B. Expected Savings: Reduced Form



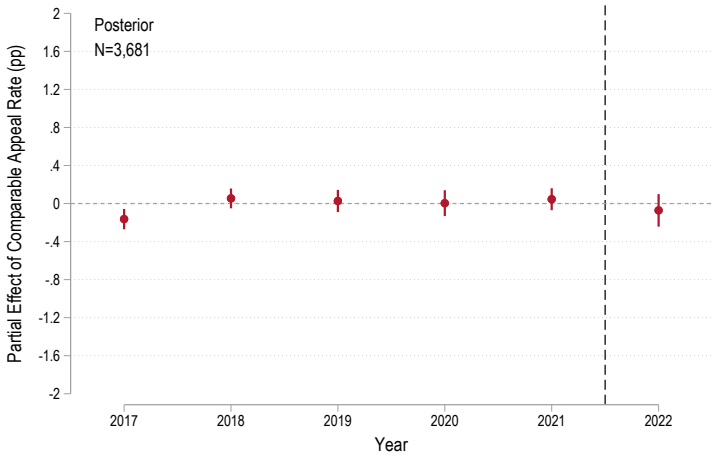
C. Richest -1% Appeal Rate: Event-Study



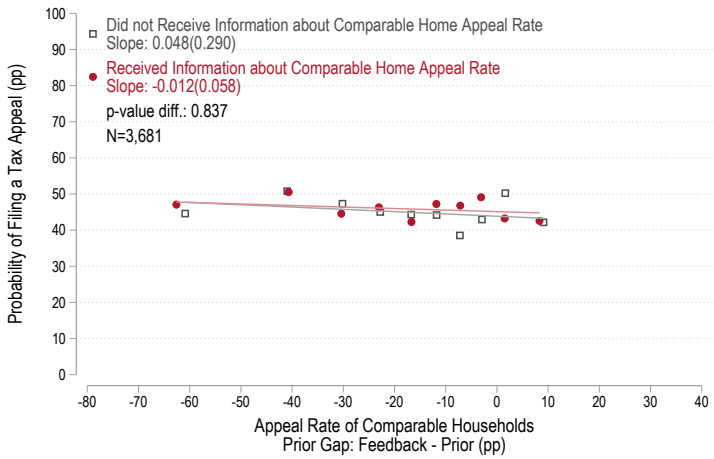
D. Richest -1% Appeal: Reduced Form



E. Comparable Appeal Rate: Event-Study



F. Comparable Appeal Rate: Reduced Form



Notes: This figure is analogous to Figure 7 in the main text but reports event-study and reduced-form binned scatterplot estimates from regressions capturing partial effects rather than total effects.

Table A.1: Characteristics of Respondents to the Field Survey

		Subject Pool		
	Letter (1)	Yes (2)	No (3)	p-value test (2 vs. 3)
2022 Home Value ($1,000s)	376.982	448.006	374.285	< 0.001
	(0.814)	(4.224)	(0.829)	
2022 Property Tax Amount ($1,000s)	8.513	9.697	8.468	< 0.001
	(0.017)	(0.088)	(0.017)	
Homestead Exemption (%)	65.304	87.014	64.480	< 0.001
	(0.150)	(0.554)	(0.154)	
Homestead Exemption and Binding Cap (%)	45.094	66.205	44.292	< 0.001
	(0.157)	(0.780)	(0.160)	
2021 Appeal Rate (%)	4.840	9.318	4.670	< 0.001
	(0.068)	(0.479)	(0.068)	
2020 Appeal Rate (%)	8.556	17.169	8.229	< 0.001
	(0.088)	(0.622)	(0.088)	
2019 Appeal Rate (%)	6.343	11.464	6.149	< 0.001
	(0.077)	(0.525)	(0.077)	
2018 Appeal Rate (%)	6.115	10.052	5.966	< 0.001
	(0.076)	(0.496)	(0.076)	
Multiple Owners (%)	24.199	30.128	23.974	< 0.001
	(0.135)	(0.756)	(0.137)	
Living Area (1,000s Sq. Feet)	2.063	2.302	2.054	< 0.001
	(0.003)	(0.014)	(0.003)	
Number of Bedrooms	3.301	3.419	3.296	< 0.001
	(0.002)	(0.012)	(0.002)	
Number of Baths	2.101	2.287	2.094	< 0.001
	(0.002)	(0.012)	(0.002)	
Observations	100,615	3,681	96,934	

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 all subjects in the field experiment who were mailed a letter. Columns (2)–(3) break down that sample into households that responded to the Field Survey and households that did not respond to the Field Survey, respectively. Column (4) reports the p-value of the test of equal means across the respondent and non-respondent groups.

Table A.2: Summary of Feedback Provided in the Information-Provision Experiment

Subject's Own Home Value	Percent of Subjects	Richest-1% Appeal Rate	Comparables Appeal Rate	Avg. Savings
\$50,000 - \$99,999	0.73%	49%	5%	\$240
\$100,000 - \$149,999	3.21%	49%	5%	\$289
\$150,000 - \$199,999	10.46%	49%	6%	\$294
\$200,000 - \$249,999	15.95%	49%	8%	\$322
\$250,000 - \$299,999	14.67%	49%	11%	\$377
\$300,000 - \$349,999	11.95%	49%	13%	\$400
\$350,000 - \$399,999	9.70%	49%	17%	\$440
\$400,000 - \$449,999	7.44%	49%	20%	\$526
\$450,000 - \$499,999	6.28%	49%	23%	\$650
\$500,000 - \$549,999	3.94%	49%	25%	\$692
\$550,000 - \$599,999	3.53%	49%	25%	\$775
\$600,000 - \$649,999	2.58%	49%	30%	\$793
\$650,000 - \$699,999	1.79%	49%	30%	\$750
\$700,000 - \$749,999	1.52%	49%	31%	\$841
\$750,000 - \$799,999	1.11%	49%	34%	\$905
\$800,000 - \$849,999	0.95%	49%	32%	\$980
\$850,000 - \$899,999	0.46%	49%	35%	\$1,104
\$900,000 - \$949,999	0.44%	49%	37%	\$1,138
\$950,000 - \$999,999	0.41%	49%	35%	\$1,304
\$1,000,000 - \$1,049,999	0.46%	49%	38%	\$1,119
\$1,050,000 - \$1,099,999	0.41%	49%	36%	\$1,091
\$1,100,000 - \$1,149,999	0.35%	49%	35%	\$1,334
\$1,150,000 - \$1,199,999	0.35%	49%	37%	\$1,405
\$1,200,000 - \$1,249,999	0.14%	49%	38%	\$1,204
\$1,250,000 - \$1,299,999	0.22%	49%	39%	\$1,139
\$1,300,000 - \$1,349,999	0.16%	49%	41%	\$1,236
\$1,350,000 - \$1,399,999	0.22%	49%	41%	\$1,343
\$1,400,000 - \$1,449,999	0.11%	49%	43%	\$1,840
\$1,450,000 - \$1,499,999	0.14%	49%	42%	\$1,561
\$1,500,000 - \$1,549,999	0.11%	49%	46%	\$1,714
\$1,550,000 - \$1,599,999	0.14%	49%	41%	\$1,890
\$1,600,000 - \$1,649,999	0.08%	49%	45%	\$1,795
\$1,650,000 - \$1,699,999	0%	49%	41%	\$1,740
\$1,700,000 - \$1,749,999	0%	49%	45%	\$2,015
\$1,750,000 - \$1,799,999	0%	49%	49%	\$1,909
\$1,800,000 - \$1,849,999	0%	49%	50%	\$1,600

Notes: The table reports the feedback provided in the information provision experiment. The statistics were computed using data from the DCAD administrative records, corresponding to the 2021 appeals. The first column lists the home value bins of the subjects. The second column shows the percentage of subjects in each bin. The third column reports the appeal rate of the richest-1%, which was the same value (49%) for all subjects. The fourth presents the correct appeal rate of comparable homes. The last column presents the correct average tax savings for comparable homes.

Table A.3: Characteristics of Respondents to Surveys

	Survey Type		
	Field (1)	Attitude (2)	p-value (3)
Republican (%)	32.723	47.667	< 0.001
	(0.794)	(2.041)	
Democrat (%)	49.170	43.333	0.008
	(0.846)	(2.025)	
Independent (%)	18.106	9.000	< 0.001
	(0.651)	(1.169)	
Female (%)	41.726	48.667	0.001
	(0.831)	(2.042)	
Age	47.005	40.207	< 0.001
	(0.188)	(0.538)	
Race: White (%)	55.436	62.833	0.001
	(0.838)	(1.975)	
Race: Black (%)	11.893	24.500	< 0.001
	(0.545)	(1.757)	
Race: Latino (%)	18.450	4.000	< 0.001
	(0.654)	(0.801)	
Race: Other (%)	14.420	2.333	< 0.001
	(0.592)	(0.617)	
Lives in Texas (%)	100.000	26.167	< 0.001
	(0.000)	(1.796)	
Observations	3,681	600	

Notes: Average pre-treatment characteristics of subjects in the field experiment, with standard errors in parentheses. Column (1) corresponds to subjects who responded to the Field Survey. Column (2) corresponds to subjects in the attitudes survey. Column (3) reports the p-value of the test of equal means across the groups.

Table A.4: Full List of Tax Behaviors included in the Attitudes Survey

Tax Behavior	**Explanation**
Property Tax Appeals	You can lower your property taxes by filing a tax appeal. For example, if the county says your home is worth $400,000, you can file an appeal to argue that it is worth $375,000. If you win the appeal—like many homeowners do—your property taxes will go down because they'll now be based on the lower $375,000 assessment. Some homeowners go through this process themselves, while others hire professionals to handle it.
Tax Deferral Through Retirement Accounts	People lower their taxable income by contributing to retirement accounts like 401(k)s and IRAs. Instead of paying taxes immediately, they delay them until they withdraw the money, usually in retirement when their tax rate may be lower. This strategy lets taxpayers legally reduce their tax burden by shifting taxes to a future time when they might owe less. Contribution limits, withdrawal timing, and tax law changes affect how well this works.
Buy, Borrow, Die	Wealthy individuals avoid taxes by borrowing against their investments instead of selling them. They take out loans using stocks or real estate as collateral, allowing them to access money without triggering capital gains taxes. When they die, their heirs inherit the assets at a new tax value, erasing the unpaid capital gains taxes. This strategy lets individuals legally reduce their tax burden by replacing taxable income with tax-free loan money. Estate planning and loan costs affect its success.
Harvesting Tax Losses	Investors reduce their tax bills by selling assets at a loss to cancel out taxes on profitable investments. If their losses exceed gains, they lower their taxable income up to a limit and carry forward extra losses to future years. This strategy lets investors legally reduce their tax burden by balancing their investment profits with losses. Market trends, tax laws, and timing play key roles in how well it works.
Real Estate Depreciation	Property owners lower their taxable income by claiming depreciation, which treats buildings as losing value over time. Even if their property gains market value, they still deduct part of its cost each year. This strategy lets real estate investors legally reduce their tax burden by reporting rental income as lower than it actually is. IRS rules, property type, and long-term tax planning determine how much they save.
Carried Interest	Investment managers classify their earnings as capital gains instead of regular income to pay a lower tax rate. Instead of paying higher income tax rates on their compensation, they pay the lower long-term capital gains tax rate. For example, a private equity manager who earns a $10 million performance bonus can structure it as carried interest, reducing their tax rate from 37% (ordinary income tax) to 20% (capital gains tax). This allows investment managers to reduce their tax liability using rules intended for investors rather than salaried workers.
Profit Shifting	Multinational companies lower their tax bills by moving profits to countries with lower tax rates. They adjust prices on sales between their own branches, charge overseas offices for brand rights, or register businesses in tax-friendly locations. This strategy lets corporations legally reduce their tax burden by reporting income in low-tax countries instead of high-tax ones. Global tax laws, corporate structure, and government regulations shape how effective this is.
Tax Avoidance through Substitution	Consumers save money by choosing lower-taxed products instead of higher-taxed ones. For example, a smoker might buy rolling tobacco instead of pre-packaged cigarettes because it's taxed less. Others switch to vaping or e-cigarettes, which some governments tax less. This strategy lets individuals legally reduce their tax burden by selecting products with lower tax rates. Local tax laws, availability, and government regulations determine how much they save.
Duty Free Stores	Travelers lower their tax costs by shopping at duty-free stores in airports, cruise ports, and border locations. These stores sell alcohol, perfume, and luxury goods without import taxes or sales tax. Since the goods are meant for international travelers, they do not face standard retail taxes. This strategy lets individuals legally reduce their tax burden by buying products in tax-exempt zones. The amount they save depends on customs limits and enforcement.
Under Reporting Income	A person or business can fail to report all of their income to tax authorities. This can include receiving cash payments and not declaring them, underreporting sales, or misclassifying income. By doing this, they avoid paying taxes that would otherwise be due.
Paying Workers under the Table	An employer can pay workers in cash without reporting the wages to tax authorities. This allows them to avoid payroll taxes, Social Security, Medicare, and unemployment insurance contributions. Employees who receive cash payments do not have taxes withheld and may not have access to benefits like Social Security, workers' compensation, or unemployment insurance.
Falsifying Deductions or Expenses	A taxpayer can claim false deductions or inflate business expenses on their tax return. This can include overstating business costs, listing personal purchases as business expenses, or fabricating dependents to receive tax credits. By reducing their taxable income, they decrease the amount of taxes owed.

Table A.5: Correlates of Willingness to Help Others Appeal

Group	**Share Favoring Aid to**		**Share Opposing Aid to**	
	Poorest-1% (1)	**Richest-1%** (2)	**Poorest-1%** (3)	**Richest-1%** (4)
All [N=3,502]	90.520	40.320	5.540	41.319
	(0.495)	(0.829)	(0.387)	(0.832)
Female [N=1,458]	91.975	38.477	4.458	42.112
	(0.712)	(1.275)	(0.541)	(1.294)
Male [N=2,032]	89.567	41.634	6.201	40.699
	(0.678)	(1.094)	(0.535)	(1.090)
Self-Identify as Democrat [N=1,704]	91.960	33.099	4.636	48.415
	(0.659)	(1.140)	(0.510)	(1.211)
Self-Identify as Republican [N=1,137]	90.062	49.868	5.981	31.135
	(0.888)	(1.483)	(0.704)	(1.374)
Self-Identify as Independent [N=622]	88.264	41.961	6.431	40.836
	(1.292)	(1.980)	(0.984)	(1.972)
Below Median Home Value [N=922]	90.347	40.130	5.206	42.408
	(0.973)	(1.615)	(0.732)	(1.628)
Above Median Home Value [N=2,580]	90.581	40.388	5.659	40.930
	(0.575)	(0.966)	(0.455)	(0.968)
Appeal Inequality Unfair [N=954]	91.405	25.996	5.136	56.080
	(0.908)	(1.421)	(0.715)	(1.608)
Appeal Inequality Fair [N=2,548]	90.188	45.683	5.691	35.793
	(0.589)	(0.987)	(0.459)	(0.950)
Owner Appeal [N=1,620]	91.728	42.284	4.506	39.321
	(0.685)	(1.228)	(0.516)	(1.214)
Owner Did Not Appeal [N=1,882]	89.479	38.629	6.429	43.039
	(0.566)	(1.142)	(0.566)	(1.142)

Notes: This table shows the share of subjects with positive (favoring) or negative (opposing) willingness to pay to help a homeowner in a given part of the house value distribution appeal their property taxes. Standard errors in parentheses. Indifference (WTP = \$0) towards aid for each group is omitted.

Table A.6: Main Results: 2SLS with Reduced-Form and First-Stage Results

	Protested in 2022	
	Total (1)	Partial (2)
A. Effects of Beliefs on Appeals (2SLS)		
Expected Savings ($100s)	1.302**	0.583***
	(0.570)	(0.206)
Richest-1% Appeal Rate	−0.034	0.126
	(0.142)	(0.086)
Comparable Appeal Rate	−0.054	−0.072
	(0.157)	(0.104)
B. Reduced Form		
$T_i^s \times (s_{i,t-1}^{\text{end}} - s_{i,t-1}^{\text{prior}}) \times H_i$	0.356	0.275
	(0.326)	(0.050)
$T_i^s \times (s_{i,t-1}^{\text{end}} - s_{i,t-1}^{\text{prior}}) \times (1 - H_i)$	0.671**	0.498***
	(0.291)	(0.192)
$T_i^w \times (w_{i,t-1}^{\text{end}} - w_{i,t-1}^{\text{prior}})$	−0.024	0.064
	(0.101)	(0.050)
$T_i^c \times (c_{i,t-1}^{\text{end}} - c_{i,t-1}^{\text{prior}})$	−0.038	−0.012
	(0.112)	(0.058)
C. First Stage: Savings Beliefs		
$T_i^s \times (s_{i,t-1}^{\text{end}} - s_{i,t-1}^{\text{prior}}) \times H_i$	0.294	0.434***
	(0.192)	(0.089)
$T_i^s \times (s_{i,t-1}^{\text{end}} - s_{i,t-1}^{\text{prior}}) \times (1 - H_i)$	0.678***	0.846***
	(0.139)	(0.142)
$T_i^w \times (w_{i,t-1}^{\text{end}} - w_{i,t-1}^{\text{prior}})$		−0.013
		(0.014)
$T_i^c \times (c_{i,t-1}^{\text{end}} - c_{i,t-1}^{\text{prior}})$		0.011
		(0.014)
D. First Stage: Richest-1% Beliefs		
$T_i^s \times (s_{i,t-1}^{\text{end}} - s_{i,t-1}^{\text{prior}}) \times H_i$		0.118
		(0.078)
$T_i^s \times (s_{i,t-1}^{\text{end}} - s_{i,t-1}^{\text{prior}}) \times (1 - H_i)$		0.146
		(0.327)
$T_i^w \times (w_{i,t-1}^{\text{end}} - w_{i,t-1}^{\text{prior}})$	0.709***	0.635***
	(0.030)	(0.018)
$T_i^c \times (c_{i,t-1}^{\text{end}} - c_{i,t-1}^{\text{prior}})$		0.202***
		(0.025)
E. First Stage: Comparable Beliefs		
$T_i^s \times (s_{i,t-1}^{\text{end}} - s_{i,t-1}^{\text{prior}}) \times H_i$		0.286
		(0.346)
$T_i^s \times (s_{i,t-1}^{\text{end}} - s_{i,t-1}^{\text{prior}}) \times (1 - H_i)$		-0.241
		(0.387)
$T_i^w \times (w_{i,t-1}^{\text{end}} - w_{i,t-1}^{\text{prior}})$		0.117***
		(0.018)
$T_i^c \times (c_{i,t-1}^{\text{end}} - c_{i,t-1}^{\text{prior}})$	0.708***	0.612***
	(0.053)	(0.261)
Mean outcome (baseline)	45.933	45.933
Observations:		
Expected Savings	867	
Richest-1% Appeal Rate	911	
Comparable Appeal Rate	889	
Full Sample		3,675

Notes: Significance levels: * p < 0.10, ** p < 0.05, *** p < 0.01. Robust standard errors in parentheses. Panel A: Each column reports 2SLS estimates from three regressions using Equation 5 for expected savings, richest-1% appeal rate, and comparable appeal rate on protesting in 2022. Panel B: Each column reports the reduced form effects of excluded instruments on protesting in 2022. Panel C: Each column reports the effects of excluded instruments on savings beliefs. Panel D: Each column reports the effects of excluded instruments on beliefs about the richest-1%'s appeal rate. Panel(E) reports the effects of excluded instruments on the appeal rates of comparable homes. Column (1) displays estimates of the total effects. Column (2) displays the partial effects.

Table A.7: Robustness Checks: Causal Effects of Beliefs (2SLS Estimates)

	Owner Protested in 2022				
	(1)	(2)	(3)	(4)	(5)
Panel A: Total Effects					
Expected Savings ($100s)	1.302**	1.380**	1.472	1.399**	1.364*
	(0.570)	(0.562)	(1.351)	(0.553)	(0.759)
Richest-1% Appeal Rate	-0.034	-0.016	-0.022	-0.025	-0.061
	(0.142)	(0.142)	(0.140)	(0.140)	(0.149)
Comparable Appeal Rate	-0.054	0.016	-0.022	-0.053	-0.061
	(0.157)	(0.142)	(0.140)	(0.149)	(0.149)
Kleibergen-Paap F-Statistic:					
Expected Savings	12.868	14.268	1.299	14.181	8.939
Richest-1% Appeal Rate	561.011	524.237	574.524	570.765	589.493
Comparable Appeal Rate	179.349	176.332	176.332	169.659	165.543
Observations:					
Expected Savings	867	862	889	890	776
Richest-1% Appeal Rate	911	906	928	920	799
Comparable Appeal Rate	889	880	906	900	776
Panel B: Partial Effects					
Expected Savings ($100s)	0.583***	0.570***	0.731**	0.638***	0.482**
	(0.206)	(0.207)	(0.361)	(0.208)	(0.216)
Richest-1% Appeal Rate	0.126	0.123	0.135	0.138*	0.123
	(0.086)	(0.087)	(0.084)	(0.085)	(0.088)
Comparable Appeal Rate	-0.072	-0.083	-0.065	-0.089	-0.090
	(0.104)	(0.105)	(0.102)	(0.103)	(0.111)
Kleibergen-Paap F-Statistic	15.106	15.161	15.215	14.615	13.441
Mean Outcome (Control)	45.933	47.315	46.279	45.863	47.500
Exclude Saving Outliers	✓	✓		✓	✓
Exclude Appeal Outliers	✓	✓	✓		✓
Exclude Inattentive					✓
Observations	3,681	3,496	3,735	3,742	3,229

Notes: Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. This table reports 2SLS estimates of equation 5. Column (1) reproduces the main results. Column (2) reports analogous results but includes control variables. Column (3) presents analogous results but includes the subjects who had outliers on savings beliefs. Column (4) presents analogous results but includes subjects with outliers on appeal rate beliefs. Column (5) presents analogous results but excludes subjects who failed the attention check.

B Expert Forecast Survey

We assess the extent to which experts anticipated the experimental results using a forecast survey administered to a sample of experts prior to presenting or sharing this paper. A copy of the full survey instrument is provided in Appendix H. The survey follows best practices outlined in DellaVigna et al. (2019). It begins by briefly describing the experiment and then asks experts to predict the effect of each piece of information on appeal behavior, in a format directly comparable to the experimental estimates. Following Giaccobasso et al. (2025), we elicit forecasts of the causal effect of beliefs rather than the intention-to-treat effects of information provision, which may be harder for respondents to assess. For each belief, we first elicit the effect qualitatively—that is, whether an increase in the belief would have a positive, negative, or no effect on the probability that the household files a tax appeal. After this qualitative question, we also elicit quantitative predictions. To reduce concerns that rounded responses might attenuate the elicited effect sizes, we ask for predicted effects of: (i) a 10 pp increase in beliefs about the appeal rate of the richest-1%; (ii) a 10 pp increase in beliefs about the appeal rate of comparable homes; and (iii) a $100 increase in expected savings. Lastly, our forecast survey also included two questions about the revealed-preference measure of the acceptability of tax appeals. More precisely, they were asked to predict the share of subjects who would be willing to provide help to a household from the richest-1% to file a tax appeal; and also asked about the share o households who would be willing to provide help to a household from the poorest-1%.

We collected expert responses by emailing 289 academics with publications on related topics. Invitations were sent via email during the second week of July 2023, and responses followed shortly thereafter. The final sample includes responses from 84 experts, composed of 95% professors, 4% researchers, and 1% in a policy position. Most respondents (81%) are from the field of economics; 74% report having conducted research on taxation, and 23% on preferences for redistribution.

We begin with the qualitative predictions about the effects of each of the three beliefs. A first indication that the results are not obvious to experts is their lack of confidence in their forecasts. Only 11% described themselves as "very confident" or "extremely confident" about the effects of the tax savings, 10% about the effects of the richest-1% appeal rate, and 10% about the effects of the comparable appeal rate. Their predictions are shown in Panels A–C of Figure B.1, one panel for each belief. Most experts (66%) forecasted that higher beliefs about tax savings would increase the likelihood of filing an appeal, which aligns with our experimental evidence. A somewhat larger share (70%) predicted that beliefs about the appeal rate of the richest-1% would raise appeal rates, a prediction not borne out in the data.

An even greater majority (76%) anticipated positive effects from beliefs about the appeal rate of comparable homes, which again contradicts our null result.

In turn, Figure B.2 presents the quantitative forecasts, scaled when appropriate to allow direct comparison with the experimental estimates.[65] Panel A plots the predicted effect of a \$100 increase in expected savings, Panel B the effect of a 1 pp increase in the richest-1% appeal rate, and Panel C the effect of a 1 pp increase in the comparable appeal rate. In each panel, the green line depicts the cumulative distribution function of forecasts, the vertical gray line marks the experimental point estimate, and the shaded area shows the corresponding 90% confidence interval. Importantly, contrary to expert expectations, perceptions of expected tax savings had a stronger effect on the decision to file a protest than beliefs about the appeal behavior of the richest-1% or comparable properties. Overall, the quantitative forecasts mirror the qualitative ones: experts tended to underestimate the effect of tax savings but overestimate the effects of perceived appeal rates for both the richest-1% and comparable households.

Lastly, Panel D of Figure B.1 reports the forecasts for the revealed-preference measure of the acceptability of tax appeals. About 9.8% of experts expected homeowners to be willing to pay to help someone in the richest-1% reduce their tax burden, while 43.9% predicted homeowners would pay to help someone in the poorest-1%. Expert forecasts were largely inaccurate in that they substantially underestimated the willingness to help both groups. This inaccuracy aligns with their low confidence, as only 6.4% of respondents described themselves as "very confident" or "extremely confident" in these forecasts. On the other hand, the predictions were directionally accurate in a *relative* sense, as experts anticipated greater acceptance for helping the poorest-1% than the richest-1%.

[65] We consider the total effects specification the most relevant benchmark, since the forecast elicitation did not instruct experts to hold other beliefs constant.

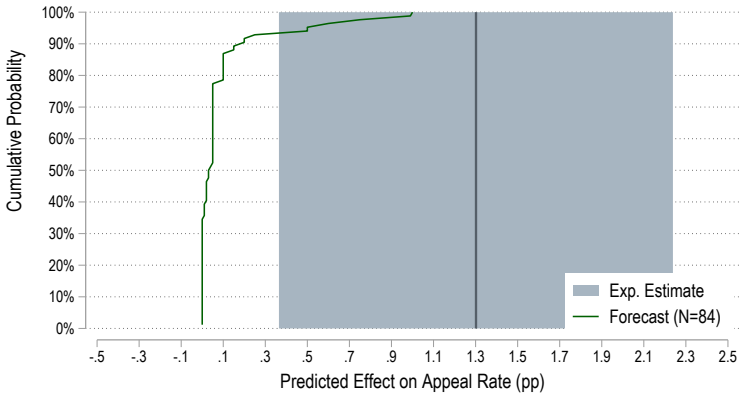
Figure B.1: Expert Forecasts



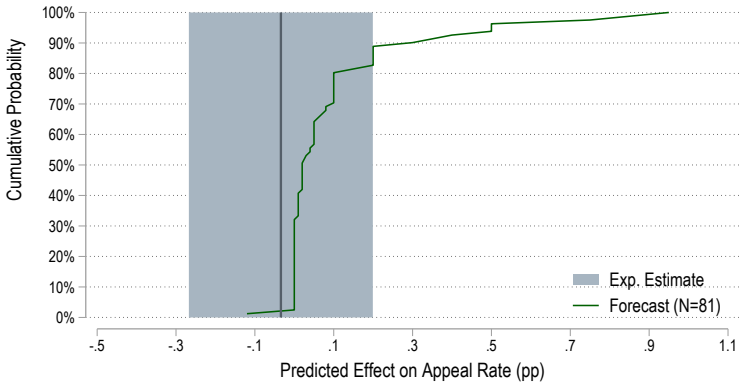
Notes: Panel A displays the results to the following question: "Suppose that at the beginning of the survey, a respondent expected that, among the richest households in the county, 50% of them will file a tax appeal in 2022. Due to the information shock, the respondent now believes that, among the richest households in the county, 60% of them will protest that is, the belief increased by 10 pp. As a result of this belief update, would you expect the respondent's probability of filing a protest to go up, stay the same or go down?" Panel B displays responses to the question: "Suppose that at the beginning of the survey, a respondent expected that, among similar households in the county, 20% of them will file a tax appeal in 2022. Due to the information shock, the respondent now believes that, among similar households in the county, 30% of them will protest that is, the belief increased by 10 pp. As a result of this belief update, would you expect the respondent's probability of filing a protest to go up, stay the same or go down?". Panel C displays responses to the question: "Suppose that at the beginning of the survey, a respondent expected to save \$800 in 2022 property taxes by filing a tax appeal. Due to the information shock, the respondent now believes that he or she can save \$900 that is, the respondent expects to save an additional \$100. As a result of this belief update, would you expect the respondent's probability of filing a protest to go up, stay the same or go down?". Panel D displays the portion of the sample that wanted to help other taxpayers protest along with responses to the questions: "What share of respondents do you think that wanted us to help one of the richest households (top-1%) (bottom-1%) to file a protest?"

Figure B.2: Comparison between Quantitative Expert Forecasts and Experimental Estimates

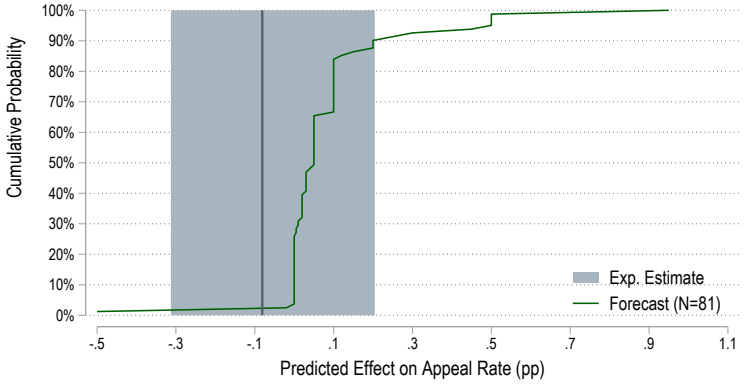
A. Effect of Expected Savings



B. Effect of Richest-1% Appeal Rate

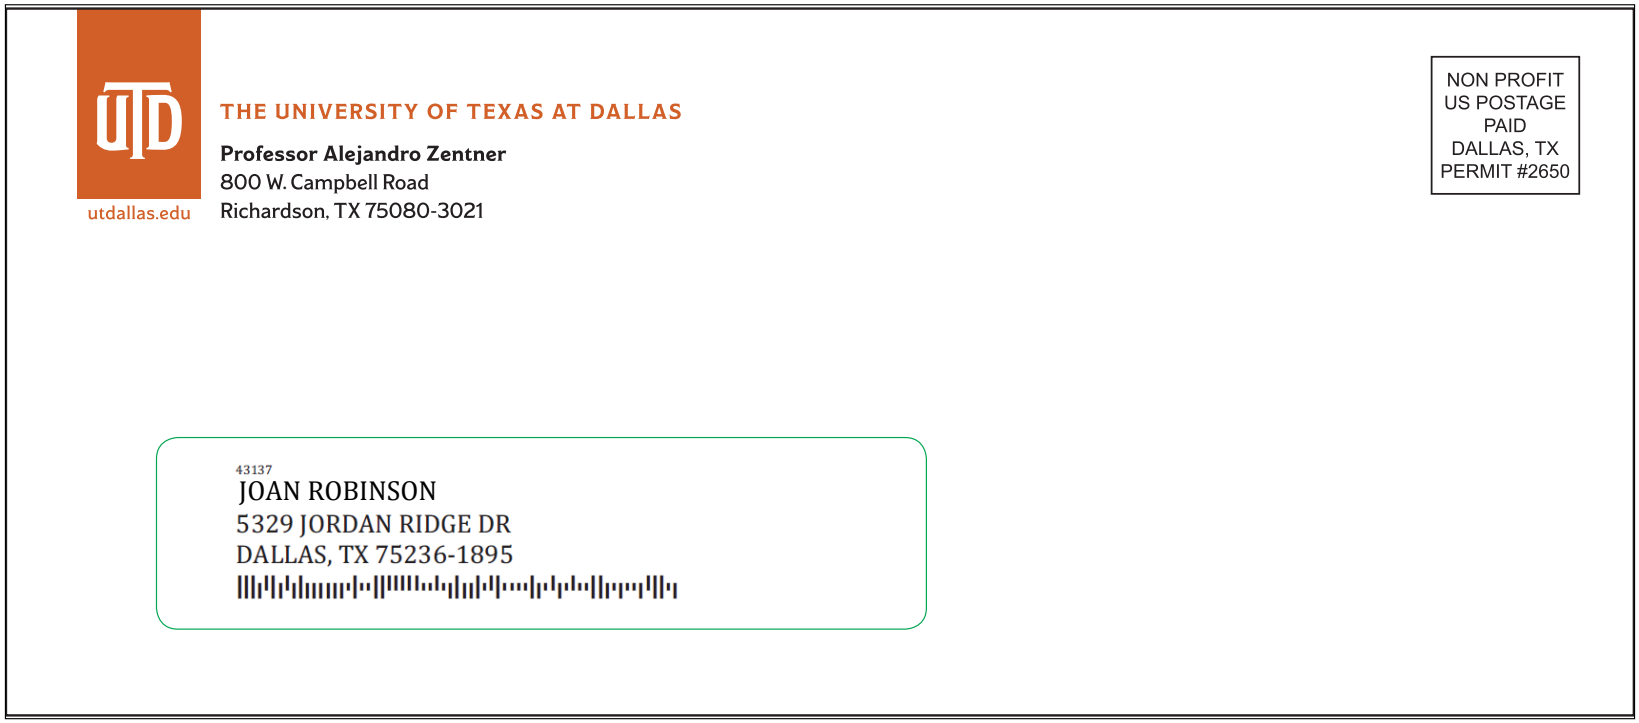


C. Effect of Comparable Appeal Rate



Notes: This figure shows the distribution of expert predictions about the effects of a 1 pp increase in beliefs on the probability that a homeowner files a protest directly. To make the elicitation easier, in the prediction survey, for the predictions reported in Panels B and C, we asked subjects to predict the effects of a 10 pp increase in beliefs. For this figure, we divide those two predictions by 10 and we obtained the effect per 1 pp so these coefficients can be compared directly to the 2SLS estimates. The vertical red solid line corresponds to the experimental estimate based on the 2SLS specification. The shaded area (in pink) corresponds to the 90% confidence interval. Panel A displays the information for expected savings. Panel B displays the information for richest-1% appeal rates. Panel C displays the information for peer appeal rates.

C Sample Envelope



D Sample of Full Letter

THE UNIVERSITY OF TEXAS AT DALLAS
Naveen Jindal School of Management

April 20th, 2022

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 about property taxes. Find below the estimated 2022 taxes for your home at 1504 Tuley St (Cedar Hill, TX) in Dallas County:

	YOUR HOME
Proposed Value	*$248,040*
Estimated Tax Amount	*$6,330*

Source: Data provided by Dallas Central Appraisal District (CAD). Proposed Value is Dallas CAD's estimate of your home's market value as of January 1st, 2022. Estimated Tax Amount is our estimate of taxes due this year using the latest tax rates available (some exemptions might not be included).

You can lower your tax amount by protesting the proposed value of your property. **If you would like to help us with our study, we kindly ask you to fill out a short survey:**

Visit https://www.utdallas.edu/taxsurvey/ and enter validation code **930661**

Please respond to the survey as soon as possible, regardless of whether you plan to file a protest or not.

It takes only a few minutes to respond to the survey, and your responses will be confidential. At the end of our survey, we provide step-by-step instructions on how to file a protest online or by mail, if you wish to do so.

The deadline to file a protest is May 16th, 2022. With our instructions, filing a protest is really simple. You do not need an agent. You do not need to attend a hearing if you accept an online settlement offered by the county. If the county schedules a hearing and you do not attend it, the protest will simply be dismissed with no penalty.

800 W Campbell Rd, Office 3206
Richardson, TX 75080

Please recycle

Your household was randomly chosen to receive this letter. If you have any questions about the study, you can find contact information on the study's website provided at the end of the survey.

To access the survey, you can also use the following QR code:



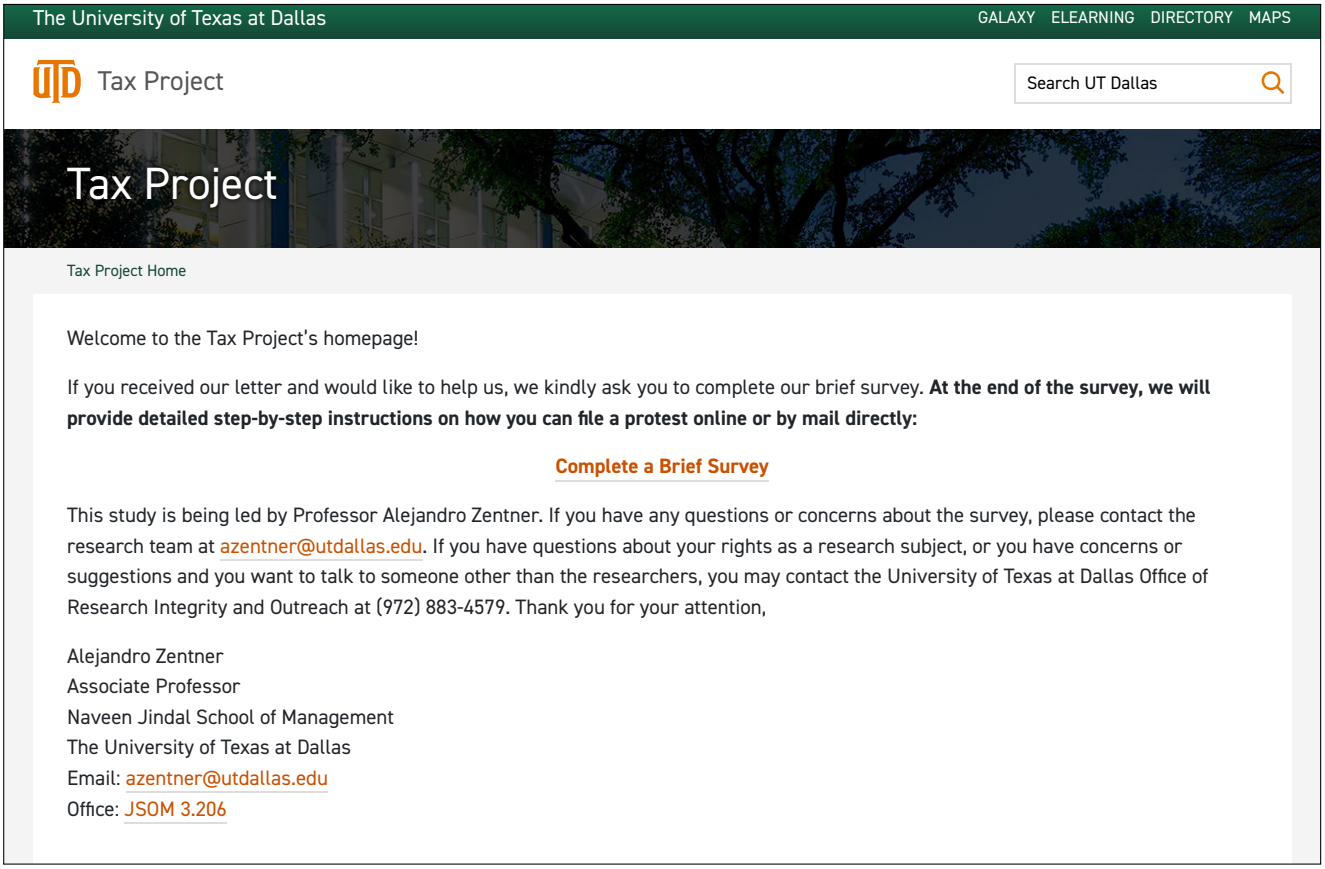
Thank you for your attention!

Alejandro Zentner

Associate Professor
The University of Texas at Dallas
https://www.utdallas.edu/taxproject/

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

E Project's Website



F Questionnaire: Field Experiment

Welcome to our survey on views about property taxes. Please read the consent form below and click "I Agree" when you are ready to start the survey.

At the end of the survey, we will provide detailed step-by-step instructions on how you can file a protest online or by mail directly. These instructions make filing a protest really simple. You will not need an agent. You will not need to attend a hearing if you accept an online settlement offered by the county. If the county schedules a hearing and you do not attend it, the protest will simply be dismissed with no penalty.

The study is being conducted by a team of researchers led by Professor Alejandro Zentner of The University of Texas at Dallas, and it has been designated by The University of Texas at Dallas Office of Research Integrity and Outreach as exempt from review by an Institutional Review Board. No deception is involved, and the study involves no more than minimal risk to participants (i.e., the level of risk encountered in daily life). Participation in the study typically takes about 10 minutes and is strictly confidential. Participants begin by entering the validation code included in the letter received by mail and then answer questions related to property taxes and demographics. All responses are treated as confidential. Data will be pooled and published in aggregated form only. Participants should be aware; however, that although the survey is being run from a "secure" https server of the kind typically used to handle credit card transactions, there is a small possibility that responses could be viewed by unauthorized third parties (e.g., computer hackers). Many individuals find participation in this study enjoyable, and no adverse reactions have been reported thus far. Participation is voluntary, refusal to take part in the study involves no penalty or loss of benefits to which participants are otherwise entitled, and participants may withdraw from the study at any time without penalty or loss of benefits to which they are otherwise entitled. By filling the survey, you may earn up to $30 in the survey. Our best guess is that the odds of winning up to $30 are 1 in 100. If participants have further questions about this study, they may contact the Principal Investigator, Alejandro Zentner (azentner@utdallas.edu). Participants who want more information about their rights as a participant or who want to report a research related concern may contact The University of Texas at Dallas Office of Research Integrity and Outreach at (972) 883-4579. If you are 18 years of age or older, understand the statements above, and freely consent to participate in the study, click on the "I Agree" button to begin the survey.

◯ I agree, I would like to take part in this study and confirm that I am 18 years of age or older, I understand the statements above, and freely consent to participate in the study.

Please log in.

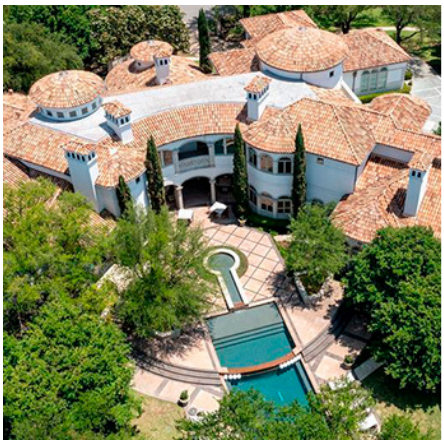
Validation Code

We want to ask you some questions about a property you own in Dallas County: 941 Trinity Dr (Lancaster, TX). In 2022, the Dallas Central Appraisal District (DCAD) assessed this property at $329,880, with estimated property taxes of $7,797.

Have you already filed a property tax protest in 2022?

- ◯ Yes
- ◯ No

Consider the richest 1% of households in Dallas County: households with homes assessed at $1.9 million or more in **2021**.





What percent of these households do you think filed a property tax protest last year (in 2021)?

0 10 20 30 40 50 60 70 80 90 100

Now, consider households from Dallas County that are similar to your household: households with homes assessed at $250,000 - $299,999 in **2021**. In comparison, your home was assessed at $282,530 in 2021.

What percent of these households do you think filed a property tax protest last year (in **2021**)?

0 10 20 30 40 50 60 70 80 90 100

Consider households from Dallas County that are similar to your household (homes assessed at $250,000 - $299,999 in 2021) **and filed a tax protest**.

Last year (in 2021), before they filed a tax protest, these households were expected to pay, on average, $5,271 in property taxes.

How much do you think these households saved in 2021 property taxes, on average, due to their tax protests?

$

Next, a group of survey respondents will be randomly chosen to receive some information related to the previous questions.

Please continue to the next screen to find out if you were selected to receive the information.

You have been randomly selected to receive the following information:

Among the richest 1% of households in Dallas County (homes assessed at $1.9 million or more in 2021), 49% filed a tax protest in 2021.

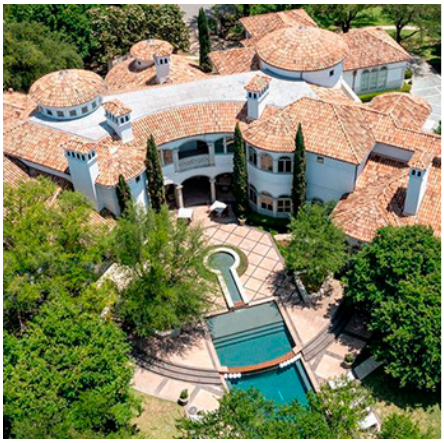
- **Among households from Dallas County that are similar to your household (homes assessed at $250,000 - $299,999 in 2021), 11% filed a tax protest in 2021.**

- **Among households from Dallas County that are similar to your household and filed a tax protest in 2021, they saved an average of $377 in 2021 property taxes.**

Please take a moment to review this information carefully, because you will not be able to return to this screen.

When ready, proceed to the next screen to continue with the survey.

Previously, we asked a few questions about property taxes during the last year (2021 taxes). **Now, we want to ask you the same questions, but about this year (2022).**





Consider the richest 1% of households in Dallas County: households with homes assessed at $1.9 million or more in **2021**.

What percent of these households do you think will file a property tax protest this year (in 2022)?

0 10 20 30 40 50 60 70 80 90 100

Now, consider households from Dallas County that are similar to your household: households with homes assessed at $250,000 - $299,999 in **2021**.

What percent of these households do you think will file a property tax protest this year (in 2022)?



In 2022, if you do not file a tax protest, we estimate that you will have to pay $7,797 in property taxes.

If you file a tax protest in 2022, how much do you expect to save in 2022 property taxes?

$

You have time until May 16th, 2022 to protest Dallas CAD's proposed value of your property. How likely are you to protest this year?

- ◯ Very likely
- ◯ Likely
- ◯ Unlikely
- ◯ Very unlikely

If you can, please explain why you will (or will not) protest in 2022:

How justifiable do you believe it is to lawfully reduce your tax bill if you have a chance?

- ◯ Always justified
- ◯ Sometimes justified
- ◯ Never justified

Homeowners can file a protest on their own, but without our instructions it can be difficult. We can send tailored information to households for free, so that they can effortlessly appeal if they want.

Next, we will present you with **up to 10** scenarios. In each scenario, you will be given the choice of either receiving extra money as part of your reward for answering the survey OR having us help a homeowner reduce their tax bill.

We will randomly choose **1 out of 100** respondents. If you are one of these selected respondents, we will randomly select one of your choices to be implemented. As a result, it is in your best interest to respond honestly to these scenarios.

These are real scenarios: there is a chance that you are selected, and one of your decisions will be implemented, so please answer carefully.

We can reach out to a property owner in Dallas County, whose home was assessed at $75,000 or less in 2021 (this is the poorest 1% of homes).





Would you like us to help the property owner reduce their tax bill by protesting their property value assessment?

- ◯ Yes
- ◯ No





Which of the following do you prefer?

- ◯ We help one property owner whose house was assessed at $75,000 or less reduce their taxes by protesting their property value assessment.
- ◯ You receive a $15 bonus.





Which of the following do you prefer?

- ○ We help one property owner whose house was assessed at $75,000 or less reduce their taxes by protesting their property value assessment.
- ○ You receive a $24 bonus.

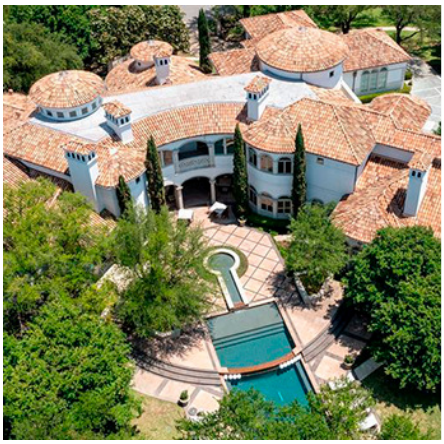




Which of the following do you prefer?

- ○ We help one property owner whose house was assessed at $75,000 or less reduce their taxes by protesting their property value assessment.
- ○ You receive a $18 bonus.

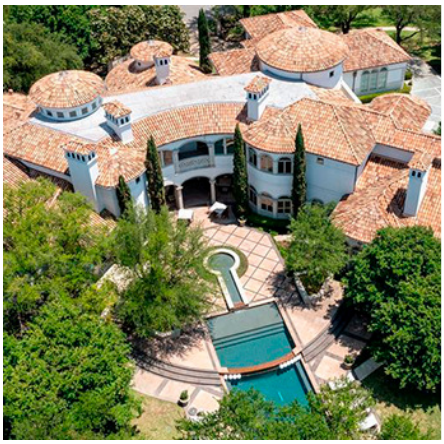
We can reach out to a property owner in Dallas County, whose home was assessed at $1.9 million or more in 2021 (this is the richest 1% of homes).





Would you like us to help the property owner reduce their tax bill by protesting their property value assessment?

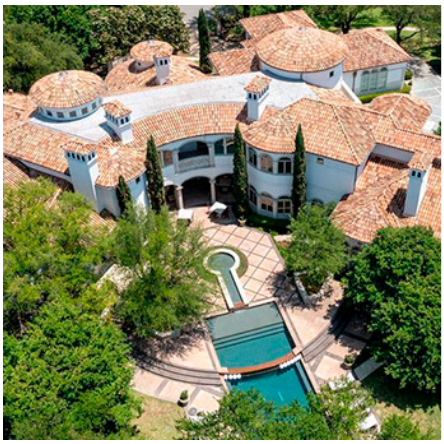
- ◯ Yes
- ◯ No





Which of the following do you prefer?

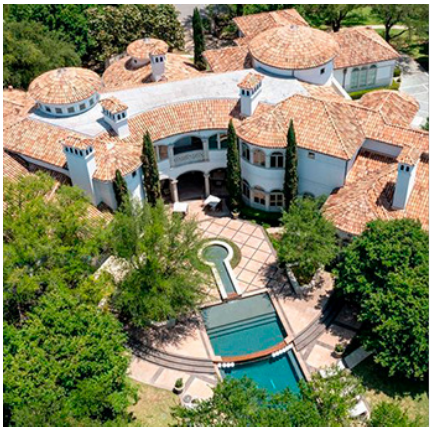
- ○ You receive a $15 bonus AND we help one property owner whose house was assessed at $1.9 million or more reduce their taxes by appealing their property value assessment.
- ○ You receive no bonus.





Which of the following do you prefer?

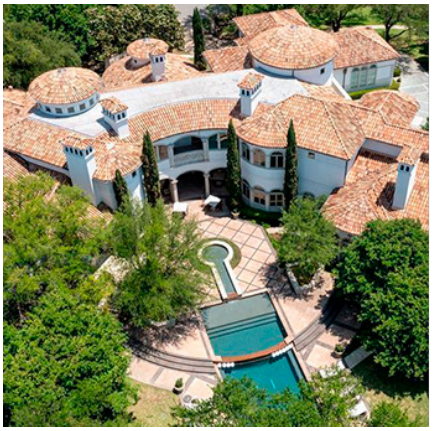
- ○ You receive a $9 bonus AND we help one property owner whose house was assessed at $1.9 million or more reduce their taxes by appealing their property value assessment.
- ○ You receive no bonus.





Which of the following do you prefer?

- ◯ You receive a $6 bonus AND we help one property owner whose house was assessed at $1.9 million or more reduce their taxes by appealing their property value assessment.
- ◯ You receive no bonus.





Which of the following do you prefer?

- ◯ You receive a $3 bonus AND we help one property owner whose house was assessed at $1.9 million or more reduce their taxes by appealing their property value assessment.
- ◯ You receive no bonus.

We are almost done. We would like to ask you a few more questions of the survey, we will provide the detailed step-by-step instructions on how you can file a protest

Please indicate your gender:

- ○ Male
- ○ Female
- ○ Other

How old are you (in years)?

[ ]

Which of the following best describes your race/ethnicity?

Please select all that apply

- ☐ American Indian or Alaska Native
- ☐ Asian / Pacific Islander
- ☐ Black or African American
- ☐ Hispanic or Latino
- ☐ White
- ☐ Other (Please specify)

[ ]

What is the highest level of education you have completed?

- ◯ Less than high school degree
- ◯ High school degree
- ◯ College degree
- ◯ Graduate degree

You said that 17% of the richest households (homes worth over $1.9 million) file their tax protest, while 19% of households like yours (homes worth $250,000 - $299,999) file a protest. Do you think this is fair or unfair?

0 - Very unfair	1	2	3	4	5 - Neither fair nor unfair	6	7	8	9	10 - Very fair
○	○	○	○	○	○	○	○	○	○	○

Next, we would like to better understand the views of Dallas County residents about their property taxes.

Do you agree or disagree with the following statement?

The local government services that I am provided (e.g., schools, roads, hospitals) justify the total amount I pay in property taxes.

0 - Strongly disagree	1	2	3	4	5 - Neither agree or disagree	6	7	8	9	10 - Strongly Agree
○	○	○	○	○	○	○	○	○	○	○

Do you think that the property taxes on the **poorest 1%** of households should go down, stay the same or go up?

- ◯ Should go down
- ◯ Stay the same
- ◯ Should go up

Do you think that the property taxes on the **richest 1%** of households should go down, stay the same or go up?

- ◯ Should go down
- ◯ Stay the same
- ◯ Should go up

In politics, as of today, do you consider yourself a Republican, a Democrat, or an Independent?

- ◯ Republican
- ◯ Democrat
- ◯ Independent

Recent research on decision making shows that choices are affected by the context in which they are made. Differences in how people feel, in their previous knowledge, experience, and in their environment can influence the choices they make. To help us understand how people make decisions, we are interested in information about you. Specifically, whether you actually take the time to read the instructions. If you don't, some results may fail to tell us very much about decision making in the real world. To help us confirm that you have read these instructions, please ignore the question about how you are feeling. Instead, only check the "none of the above" option. Thank you very much.

☐ Interested	☐ Hostile	☐ Nervous
☐ Distressed	☐ Enthusiastic	☐ Determined
☐ Excited	☐ Proud	☐ Attentive
☐ Upset	☐ Irritable	☐ Jittery
☐ Strong	☐ Alert	☐ Active
☐ Scared	☐ Inspired	☐ None of the above

What of the following options do you agree with the most?

- ◯ I understood all of the questions in the survey.
- ◯ I understood almost all of the questions in the survey.
- ◯ I understood most of the questions in the survey.
- ◯ I understood a few questions in the survey.

You have reached the end of the survey. Thanks for your participation!

To file a protest in Dallas County, you can fill out a short form online or mail it in. You do not need an agent to protest. You do not need to attend a hearing if you accept an online settlement offered by the county. If the county schedules a hearing and you do not attend it, the protest will simply be dismissed with no penalty.

You can find **step-by-step instructions for filing a protest** at the following website: **https://taxhelp.utdallas.edu/**

Please be sure to print or save this URL (https://taxhelp.utdallas.edu/), because you may not be able to return to this page.

G Questionnaire: Attitudes Survey



Welcome to our public opinion survey. This survey is intended for homeowners only.

This survey takes approximately 5 minutes, and we anticipate no risks or discomfort from your participation.

Please click "I Agree" when you're ready to begin the survey.

○ I Agree

Click this to ensure you are not a robot





What is your Prolific ID?

Please note this response should auto-fill with the correct ID.



Taxes are essential for funding public services, yet many individuals and businesses actively seek ways to minimize their tax bills. In this survey, we'd like to understand your opinions on five different forms of tax compliance.

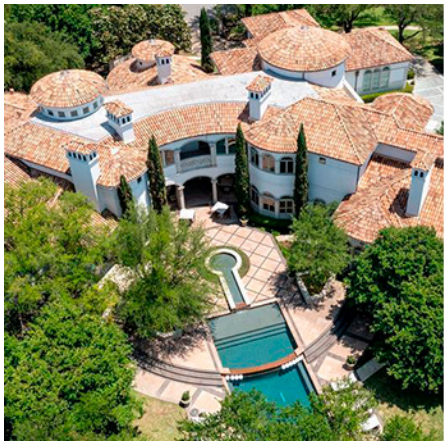
We will ask you questions about how justifiable different tax behaviors are. You'll be asked to consider this both in general and in relation to two specific groups.

First, we'll ask about the poorest 1% of households—families living in the least expensive homes, like those shown below:





Then, we'll ask about the richest 1% of households—families living in the most expensive homes, like those shown below:







Property Tax Assessment Appeals

You can lower your property taxes by filing a tax appeal. For example, if the county says your home is worth $400,000, you can file an appeal to argue that it is worth $375,000. If you win the appeal—like many homeowners do—your property taxes will go down because they'll now be based on the lower $375,000 assessment. Some homeowners go through this process themselves, while others hire professionals to handle it.

How acceptable do you believe it is to use **property tax appeals** to lower your taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable

How acceptable do you believe it is for the **poorest 1%** of households to use **property tax appeals** to lower their taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable

How acceptable do you believe it is for the **richest 1%** of households to use **property tax appeals** to lower their taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable



Buy, Borrow, Die

Wealthy individuals avoid taxes by borrowing against their investments instead of selling them. They take out loans using stocks or real estate as collateral, allowing them to access money without triggering capital gains taxes. When they die, their heirs inherit the assets at a new tax value, erasing the unpaid capital gains taxes. This strategy lets individuals legally reduce their tax burden by replacing taxable income with tax-free loan money. Estate planning and loan costs affect its success.

How acceptable do you believe it is to use the **"buy, borrow, die" strategy** to lower your taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable

How acceptable do you believe it is for the **poorest 1%** of households to use **the "buy, borrow, die" strategy** to lower their taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable

How acceptable do you believe it is for the **richest 1%** of households to use **the "buy, borrow, die" strategy** to lower their taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable



Duty Free Stores

Travelers lower their tax costs by shopping at duty-free stores in airports, cruise ports, and border locations. These stores sell alcohol, perfume, and luxury goods without import taxes or sales tax. Since the goods are meant for international travelers, they do not face standard retail taxes. This strategy lets individuals legally reduce their tax burden by buying products in tax-exempt zones. The amount they save depends on customs limits and enforcement.

How acceptable do you believe it is to shop at **duty free stores** to lower your taxes?

- ○ Completely unacceptable
- ○ Mostly unacceptable
- ○ Mostly acceptable
- ○ Completely acceptable

How acceptable do you believe it is for the **poorest 1%** of households to shop at **duty free stores** to lower their taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable

How acceptable do you believe it is for the **richest 1%** of households to shop at **duty free stores** to lower their taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable



Under Reporting Income

A person or business can fail to report all of their income to tax authorities. This can include receiving cash payments and not declaring them, underreporting sales, or misclassifying income. By doing this, they avoid paying taxes that would otherwise be due.

How acceptable do you believe it is to **underreport income** to lower your taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable

How acceptable do you believe it is for the **poorest 1%** of households to **underreport income** to lower their taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable

How acceptable do you believe it is for the **richest 1%** of households to **underreport income** to lower their taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable



Tax Avoidance through Substitution

Consumers save money by choosing lower-taxed products instead of higher-taxed ones. For example, a smoker might buy rolling tobacco instead of pre-packaged cigarettes because it's taxed less. Others switch to vaping or e-cigarettes, which some governments tax less. This strategy lets individuals legally reduce their tax burden by selecting products with lower tax rates. Local tax laws, availability, and government regulations determine how much they save.

How acceptable do you believe it is to use **substitution** to lower your taxes?

- ○ Completely unacceptable
- ○ Mostly unacceptable
- ○ Mostly acceptable
- ○ Completely acceptable

How acceptable do you believe it is for the **poorest 1%** of households to use **substitution** to lower their taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable

How acceptable do you believe it is for the **richest 1%** of households to use **substitution** to lower their taxes?

- ◯ Completely unacceptable
- ◯ Mostly unacceptable
- ◯ Mostly acceptable
- ◯ Completely acceptable



Before taking this survey, were you aware of any of these tax reduction strategies?

Please select all that apply.

- [ ] Property tax appeals
- [ ] Tax Deferral Through Retirement Accounts
- [ ] The "buy, borrow, die" strategy
- [ ] Harvesting Tax Losses
- [ ] Real Estate Depreciation
- [ ] Carried Interest
- [ ] Corporate Profit Shifting
- [ ] Tax Avoidance Through Substitution
- [ ] Shopping at duty free stores
- [ ] Under reporting income
- [ ] Paying Workers Under the Table
- [ ] Falsifying Deductions or Expenses



Have you ever used any of these strategies to lower your taxes?

Please select all that apply.

- [ ] Property tax appeals
- [ ] Tax Deferral Through Retirement Accounts
- [ ] The "buy, borrow, die" strategy
- [ ] Harvesting Tax Losses
- [ ] Real Estate Depreciation
- [ ] Carried Interest
- [ ] Corporate Profit Shifting
- [ ] Tax Avoidance Through Substitution
- [ ] Shopping at duty free stores
- [ ] Under reporting income
- [ ] Paying Workers Under the Table
- [ ] Falsifying Deductions or Expenses



What is the state and county of your primary residence (the place where you usually live)?

State

County

How old are you? (in years)

Which of the following best describes your ethnicity?

- White
- Black or African American
- Asian or Native Hawaiian and other Pacific Islander
- American Indian or Alaska Native
- Hispanic or Latino origin
- Other (Please specify)

Please indicate your gender:

- ◯ Male
- ◯ Female
- ◯ Other

What was your income in 2024, before taxes?

- ◯ I did not earn income in 2024
- ◯ $1 to $19,999
- ◯ $20,000 to $39,999
- ◯ $40,000 to $59,000
- ◯ $60,000 to $79,999
- ◯ $80,000 to $99,999
- ◯ $100,000 to $149,999
- ◯ $150,000 or more



In politics, as of today, do you consider yourself a Republican, a Democrat, or an independent?

- ◯ Democrat
- ◯ Republican
- ◯ Independent



If you had to pick between one of the two, which better describes you?

- ◯ Democrat
- ◯ Republican
- ◯ I would rather not say



Recent research on decision making shows that choices are affected by the context in which they are made. Differences in how people feel, in their previous knowledge, experience, and in their environment can influence the choices they make. To help us understand how people make decisions, we are interested in information about you. Specifically, whether you actually take the time to read the instructions. If you don't, some results may fail to tell us very much about decision making in the real world. To help us confirm that you have read these instructions, please ignore the question about how you are feeling. Instead, only check the "none of the above" option. Thank you very much.

☐ Interested	☐ Hostile	☐ Nervous
☐ Distressed	☐ Enthusiastic	☐ Determined
☐ Excited	☐ Proud	☐ Attentive
☐ Upset	☐ Irritable	☐ Jittery
☐ Strong	☐ Alert	☐ Active
☐ Scared	☐ Inspired	☐ None of the above

In your opinion, were the questions included in this survey easy or difficult to understand?

- ◯ Easy to understand
- ◯ Neither easy nor difficult
- ◯ Difficult to understand



Feel free to share any comments with us below. For example, let us know if there is a question you did not understand.

H Questionnaire: Forecast Survey



Qualtrics

Welcome to our prediction survey! We invite you to forecast the results of an experiment about property taxes. Please read the consent form below and click "I Agree" when you are ready to start the survey.

This survey involves no more than minimal risk to participants (i.e., the level of risk encountered in daily life). Participation typically takes between 5 and 10 minutes and is strictly confidential. Many individuals find participation in this survey enjoyable, and no adverse reactions have been reported thus far. Participation is voluntary, and participants may withdraw from the survey at any time.

Yes, I would like to take the survey

→

Introduction:

We study a sample of homeowners living in Dallas County (Texas) who have the opportunity to file a tax protest (a.k.a. tax appeal) of their property taxes, which can (legally) reduce the amount they have to pay in property taxes. We conducted a field experiment to explore some potential reasons why homeowners may choose to protest their taxes or not.

←

→

Subject Pool and Timing:

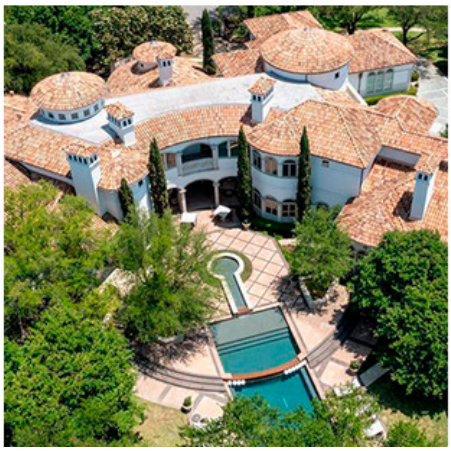
We invited a sample of homeowners from Dallas County via mail to take part in an online survey. Around 3,500 households completed the survey, which constitutes our subject pool. We conducted the survey shortly before the deadline to file a tax protest (May 16th, 2022).

Property Taxes:

Property taxes are used to fund various services such as schools, parks, and roads. The average subject had a home worth around $450,000 and was estimated to pay $10,000 in 2022 property taxes.

Tax Protests: Households have the opportunity to file a protest, by mail or online, free of charge. In their protest form households typically argue that the market value estimated by the Dallas Central Appraisal District was too high. Hereon after, we focus on whether households file a tax protest directly (i.e., without a tax agent representing them). Among households in the control group, around 45% filed a protest in 2022. Of those, 75% were successful (i.e., resulted in at least some reduction in their assessed value).

Definitions: "Similar households" are other households in the county with an assessed home value close to the corresponding assessed home value of the respondent. "Richest households" are those in the top 1% of the distribution of assessed home values in the county. These are pictures of some homes belonging to this group, which we provided to the respondents as reference:





←

→



Qualtrics

Experimental Design:

Our survey revolves around the following three beliefs:

(1) The dollar amount that the respondent expects to save in his or her 2022 property taxes if he or she filed a tax appeal this year.

(2) The share of similar households that the respondent expects to file a tax appeal this year.

(3) The share of the richest households that the respondent expects to file a tax appeal this year.

After eliciting prior beliefs, we conduct an information-provision experiment: a random sample of the subjects receive accurate information related to one (or more) of the variables listed above. This information experiment creates exogenous variation in the respondents' posterior beliefs. We use that variation to measure the causal effect of the respondents' beliefs on their subsequent behavior.

←

→

Outcome of interest:

We are interested to hear your predictions about the effects of the intervention on the main outcome of interest: whether the household filed a tax protest in 2022 (i.e., a few weeks after they completed the survey). We observe whether the subject filed a tax protest using the county's administrative records. For reference, among households in the control group (i.e., survey respondents who did not receive any information), 45% filed a tax protest in 2022.

We want to elicit your predictions about the effects of each of the three beliefs listed before. Even if you do not feel confident about it, please provide your best guess. We will ask you whether you feel confident in your guess, and also give you the opportunity to explain your reasoning (if you want).

←

→

Suppose that at the beginning of the survey, a respondent expected to save $800 in 2022 property taxes by filing a tax appeal.

Due to the information shock, the respondent now believes that he or she can save $900 – that is, the respondent expects to save an additional $100.

As a result of this update in beliefs, would you expect the respondent's probability of filing a protest to go up, stay the same or go down?

Go up

Stay the same

Go down

←

→

By how much would you expect the respondent's probability of filing a protest to go down (in percentage points)? Please enter a number between 0 and 100.

[ ] p.p.

Reference: among households in the control group, 45% filed a tax protest in 2022

How confident are you in your prediction?

- Not confident at all
- Slightly confident
- Somewhat confident
- Very confident
- Extremely confident

(Optional) Why do you expect the respondent's probability of filing a protest to go down?

Now, we want to elicit your predictions for the second belief.

Suppose that at the beginning of the survey, a respondent expected that, among similar households in the county, 20% of them will file a tax appeal in 2022. Due to the information shock, the respondent now believes that, among similar households in the county, 30% of them will protest – that is, the belief increased by 10 percentage points.

As a result of this update in beliefs, would you expect the respondent's probability of filing a protest to go up, stay the same or go down?

Go up

Stay the same

Go down

←

→

By how much would you expect the respondent's probability of filing a protest to go up (in percentage points)? Please enter a number between 0 and 100.

[ ] p.p.

Reference: among households in the control group, 45% filed a tax protest in 2022

How confident are you in your prediction?

- Not confident at all
- Slightly confident
- Somewhat confident
- Very confident
- Extremely confident

(Optional) Why do you expect the respondent's probability of filing a protest to go up?

Now, we want to elicit your predictions for the third belief.

Suppose that at the beginning of the survey, a respondent expected that, among the richest households in the county, 50% of them will file a tax appeal in 2022. Due to the information shock, the respondent now believes that, among the richest households in the county, 60% of them will protest – that is, the belief increased by 10 percentage points.

As a result of this update in beliefs, would you expect the respondent's probability of filing a protest to go up, stay the same or go down?

Go up

Stay the same

Go down

←

→

By how much would you expect the respondent's probability of filing a protest to go down (in percentage points)? Please enter a number between 0 and 100.

[ ] p.p.

Reference: among households in the control group, 45% filed a tax protest in 2022

How confident are you in your prediction?

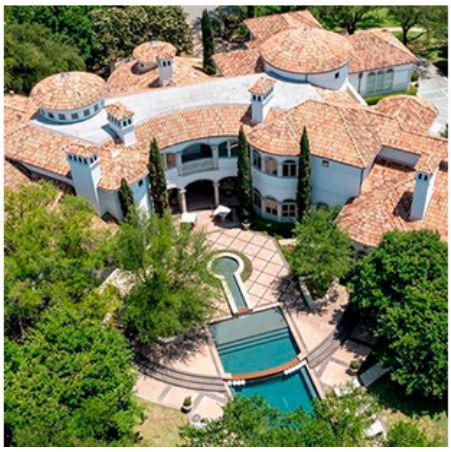
- Not confident at all
- Slightly confident
- Somewhat confident
- Very confident
- Extremely confident

(Optional) Why do you expect the respondent's probability of filing a protest to go down?

As the final part of our field experiment, we offered to help another household (different from the respondent's household) to file a protest.

Some homeowners can find it difficult to file a protest on their own. We explained to respondents that we (as researchers) could send tailored information to households for free, so that they can effortlessly appeal if they want.

First, we offered to help one household from the richest (top-1%) households. For reference, these are pictures of some of those homes, as shown to the respondents:





What share of respondents do you think wanted us to help one of the richest households (top-1%) to file a protest? (Please enter a number between 0 and 100)

[ ] %

←

→

Next, we offered to help one household from the poorest (bottom-1%) households. For reference, these are pictures of some of those homes, as shown to the respondents:





What share of respondents do you think wanted us to help one of the poorest households (bottom-1%) to file a protest? (Please enter a number between 0 and 100)

[ ] %

How confident are you in these last predictions?

- Not confident at all
- Slightly confident
- Somewhat confident
- Very confident
- Extremely confident



Qualtrics

This is the last section of the survey. We would appreciate if you could share some information about yourself.
Which of the following best describes your current position?

- Professor
- Assistant Professor
- Post-Doc
- Researcher
- PhD Student
- Master Student
- Other

Please select your discipline

- Economic
- Business (management, accounting, finance, etc.)
- Political Science
- Psychology
- Sociology
- Other

Have you ever conducted research on preferences for redistribution?

Yes

No

Have you ever conducted research on taxation?

Yes

No

←

→



Qualtrics

This is the end of the survey. If you click next, you will submit your responses.

We thank you for taking the time to provide your forecasts!

←

→