NBER WORKING PAPER SERIES

BILLIONAIRE SUPERSTAR: PUBLIC IMAGE AND DEMAND FOR TAXATION

Ricardo Perez-Truglia
Jeffrey Yusof

Working Paper 32712
http://www.nber.org/papers/w32712

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2024, Revised September 2026

This draft: September 22, 2026. First draft: July 1, 2024. We are thankful for comments by colleagues and seminar participants. We are grateful for funding provided by the Stone Center at UC Berkeley. Miriam Malament, Romina Quagliotti, Tianyi (Chelsea) Che and Alexia Witthaus Vine provided excellent research assistance. The experiment was preregistered at the AEA RCT Registry (AEARCTR-0011845) and was reviewed and approved in advance by the Institutional Review Board at UC Berkeley. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Ricardo Perez-Truglia and Jeffrey Yusof. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Billionaire Superstar: Public Image and Demand for Taxation
Ricardo Perez-Truglia and Jeffrey Yusof
NBER Working Paper No. 32712
July 2024, Revised September 2026
JEL No. D31, H24

ABSTRACT

In the United States, there are 741 billionaires with a combined net worth of $5.2 trillion. Despite this extreme inequality, billionaires face effective tax rates lower than those of the average American. We examine how the public image of billionaires shapes support for taxing them, using a pre-registered survey experiment with 9,013 Americans. Subjects were randomly assigned to one of four treatments, each emphasizing a different aspect of billionaires' public image: the role of merit in building the billionaires' wealth, their lavish consumption, the scale of their fortunes, and their average tax rates. All four treatments significantly and strongly affected the intended beliefs — for instance, information about how the billionaires made their fortunes significantly reduced the perceived role of merit. Moreover, all treatments elicited strong emotional reactions. A panel of academic experts predicted that all four treatments would substantially increase demand for taxation. In sharp contrast, we find that only one treatment—providing information about lavish lifestyles—had a robust positive effect on support for taxation, while another—focused on effective tax rates—had a strong negative effect. We discuss implications for both researchers and policymakers.

Ricardo Perez-Truglia
University of California, Los Angeles
and NBER
ricardotruglia@gmail.com

Jeffrey Yusof
University of Stuttgart
jeffrey.yusof@ivr.uni-stuttgart.de

A data appendix is available at http://www.nber.org/data-appendix/w32712
A randomized controlled trials registry entry is available at
https://www.socialscienceregistry.org/trials/11845

1 Introduction

Due to scalable markets and winner-takes-all dynamics (Rosen, 1981), a select group of individuals can accumulate immense fortunes. The billionaire class represents the epitome of this superstar phenomenon: in the United States, there are 741 billionaires with a combined net worth of $5.2 trillion.[1] Their fortunes are part of a broader trend of increasing wealth concentration at the very top (Saez and Zucman, 2020). Billionaires are also unusually visible public figures, and their extreme wealth gives them substantial opportunities to influence economic and political outcomes through campaign spending, lobbying, media ownership, and other channels. At the same time, billionaires are estimated to pay lower effective tax rates than the average American (e.g., New York Times, 2019). This combination of immense fortunes and relatively low effective tax rates has made billionaire taxation a salient policy issue, prompting several recent proposals to increase taxes on the ultra-wealthy (e.g., CBS, 2019; Tax Foundation, 2023). Despite the significance of these proposals, we know surprisingly little about the sources of public demand for such policies. This paper examines what drives that demand.

In this paper, we conduct a pre-registered survey experiment to study Americans' demand for taxing billionaires and the companies they founded. We measure the public's perceptions of billionaires, including their lifestyles, earnings, and business acumen. Furthermore, we use experimental variation to assess whether providing information about billionaires significantly affects the demand for taxation.

At the beginning of the survey, each subject is randomly assigned to one of the following five billionaires, chosen from Forbes World's Billionaires List of 2023: Elon Musk (with a net worth of approximately $180 billion at the time when our experiment was conducted), Jeff Bezos ($114 billion), Bill Gates ($111 billion), Mark Zuckerberg ($106 billion), and Michael Bloomberg ($96 billion). Our experimental design deliberately focuses on specific billionaires for several reasons—for example, these prominent individuals are among the most visible and are therefore likely to shape public perceptions of the billionaire class. Subjects are then randomized into different treatment arms. We begin by asking respondents a few questions related to the billionaire they have been assigned. Then, there is an information-provision stage, in which the respondents are randomly allocated to a piece of information about the billionaire. For example, in the luxury treatment arm, the respondent may be shown a picture of a lavish home owned by the billionaire. We measured some beliefs related to the information, both before the information-provision stage (referred to as prior beliefs) and after (referred to as posterior beliefs). These data allow us to document whether the subjects

[1] Estimates from Americans for Tax Fairness as of November 2023: https://americansfortaxfairness.org/u-s-billionaires-now-worth-record-5-2-trillion/.

had significant misperceptions initially and whether they update their beliefs in response to treatments. The last part of the survey consists of a battery of questions to serve as outcome variables, such as the respondent's preferred income tax rate for billionaires.

We study four different treatment arms, some inspired by prior research on the demand for redistribution and by the arguments made by academics, journalists, and the general public in favor of taxing the ultra-wealthy.

In the *luxury treatment*, we explore perceptions about the consumption habits of billionaires. Some billionaires seem to have made their frugality a central tenet of their public personas. For example, Warren Buffett is renowned for his modest living; he resides in the same house he purchased in the late 1950s. However, other billionaires live in extravagant homes that reflect their immense wealth. In this treatment, we provide a picture of one of the luxurious homes purchased by the billionaire, as well as information about the home's price and characteristics.

In the *non-merit treatment*, we examine perceptions about the source of a billionaire's wealth. In this arm, we randomized subjects to receive information demonstrating that factors other than the billionaire's honest and hard work played a key role in how he made his fortune. For example, we provide an excerpt from an interview with Jeff Bezos in which he explains that luck played a major role in his success with Amazon.

In the *earnings treatment*, we teach subjects about how much billionaires actually earn. We provide a side-by-side comparison between the earnings of billionaires and the earnings of the average American. For example, we mention that since 2012, Elon Musk has made $16.18 billion per year on average. Since subjects may not be able to wrap their heads around how large a billion is, we translated the information into an hourly basis: divided by the 8,760 hours in a year, Elon earned $1,847,000 per hour during this period.

In the last treatment arm, we explore perceptions about the taxes paid by billionaires. If individuals learn that billionaires pay lower tax rates than the typical American, they may conclude that billionaires are not contributing their fair share. In the *tax rate treatment*, we share the results from a report by ProPublica (2023) according to which billionaires paid an average total tax rate of 16%; in comparison, a typical American pays a rate of 21%. In addition to whether they pay low or high tax rates, the public may care about whether billionaires are playing by the rules or cheating the system. To test this hypothesis, we cross-randomized an additional piece of information. In the *tax loophole treatment*, in addition to information about tax rates, subjects receive information on two accounting strategies often used by billionaires to reduce their tax burden: receiving compensation in stocks and borrowing against them, and setting up shell companies in tax havens.

The first and main outcome of interest is the billionaire income tax. Inspired by an

actual policy proposal, we ask respondents to imagine that the U.S. government is planning to introduce a new personal income tax rate specifically for individuals earning more than $10 million annually. Respondents are given the authority to set this federal top income tax rate and are asked to use a slider to select a rate between 0% and 100%. To study whether perceptions about billionaires affect demand for taxation of the companies they founded, we describe a hypothetical corporate tax rate for large businesses and ask subjects to select a rate between 0% and 100%. Individuals may want to tax billionaires for efficiency reasons or for fairness reasons. To probe the fairness channel more directly, we also include a question on whether, from a perspective of fairness, the taxes paid by the billionaire in question are too high or too low.

In addition to the tax preferences described above, we also elicit the respondents' support for existing policy proposals. We construct an index of policy support based on six distinct measures. As part of this, respondents are asked whether they support or oppose four policy proposals, including the "Billionaire Minimum Income Tax." We let individuals split a budget in donations to two organizations, one of which has the goal of requiring big corporations and the wealthy to pay their fair share in taxes. Lastly, we give respondents the opportunity to sign a petition organized by Oxfam to increase the taxes on the ultra-wealthy. To provide complementary evidence on the mechanisms at play, we construct a non-tax attitudes index based on five different metrics such as whether the billionaire deserves his wealth or whether he is trustworthy.

To measure the persistence of treatment effects, we invited all subjects to complete an obfuscated follow-up survey about a month later. This survey again measured their posterior beliefs, along with the same outcomes measured in the baseline survey. We conducted the survey with a sample of 9,013 Americans recruited via Prolific and following best practices. Approximately 82% of these subjects also completed the follow-up survey.

We conducted a couple of complementary surveys. First, we conducted a supplemental survey, with a different subject pool, to explore individuals' emotional reactions to the treatments. Respondents were shown one of the information treatments and were asked about their reactions through a combination of multiple-choice and open-ended questions. And to determine whether the results were surprising or predictable, we conducted a forecast survey with a sample of 81 experts, primarily professors with research experience on related topics. Experts were shown each of the treatments and asked to predict its effects on the main outcome of interest: the preferred top income tax rate. The complete design and results of the expert forecast survey are presented in Appendix C.1. There is a strong consensus among experts that all treatments would have a significant and positive effect on the preferred income tax rate, with the average predicted effect ranging from around 5 to 10 percentage points

(pp), depending on the treatment.

Turning to the experimental results, we first examine whether the treatments shifted their targeted beliefs. Several of the information treatments included factual content, such as the value of a billionaire's home. By comparing respondents' prior beliefs with this information, we find that the general public holds significant misperceptions about billionaires. Most importantly, all four treatments had strong and significant effects on the targeted beliefs—for example, information about how billionaires accumulated their wealth significantly reduced the perceived role of merit. Additionally, our supplemental survey shows that each treatment elicited emotional reactions, with negative emotions significantly outweighing positive ones. Despite the significant effects on perceptions and emotional responses—and contrary to expert predictions that all four treatments would substantially increase preferred tax rates—we find that only one treatment had a robust positive effect on demand for taxation, and one treatment had strong negative effects.

The luxury treatment not only raises the average perceived value of the billionaire's home but also, more importantly, generates positive effects on the demand for taxation across the board. Relative to the control group, individuals who received the information prefer a top income tax rate that is 2.0 pp higher. The treatment increases the preferred corporate tax rate by 1.8 pp, so the effects spill over to the companies that the billionaire created. The treatment increases the perception that billionaire taxes are unfairly low by 0.08 standard deviations. And there is a significant effect on the support for tax policies of 0.08 standard deviations. A deeper analysis suggests that the effects are not driven by learning the precise value of the billionaire's home but are a product of the qualitative information, such as the narrative and the picture of the home.

The non-merit treatment leads individuals to view honest effort and hard work as less central to the billionaire's wealth. It also sharply reduces positive feelings toward the billionaire, such as admiration. By contrast, the treatment has little to no effect on support for higher taxation: the estimates are close to zero and statistically insignificant across the board. This stands in contrast with experts' predictions that the non-merit treatment would substantially raise demand for taxation. It is also puzzling given a body of experimental evidence showing that people favor more redistribution when outcomes are determined by chance rather than effort. Our preferred interpretation is that, unlike in stylized laboratory experiments, real-world economic success is perceived as the outcome of a complex interplay of factors such as skill, effort, opportunity, and even cheating. As long as legitimate factors—such as skill and effort—are perceived as playing some role, however small, it becomes difficult to justify higher taxation. Indeed, this interpretation aligns with laboratory evidence that people are more willing to accept inequalities arising from market forces or unequal circumstances, even

when those forces are outside an individual's control (Yusof ⓡ Sartor, 2024; Andre, 2024).

The earnings treatment increases subjects' perceptions of the billionaire's income. This treatment has positive effects, notably by increasing the desired income and corporate tax rates (1.9 pp and 1.46 pp, respectively). However, these effects should be taken with a grain of salt for the following reasons. First, they dissipated completely at follow-up. Second, we do not find significant effects along the other margins—the perception that billionaires pay unfairly low taxes and the support for related policies. Like the luxury treatment, the earnings treatment is intended to highlight the extreme inequalities between billionaires and the average American. However, the luxury treatment has a more robust effect on preferences for taxation. Our preferred interpretation is that images are more effective than statistics because they are easier for people to relate to.

The tax rate treatment lowers perceived tax rates faced by billionaires by about 5.5 pp on average. It also increases the perception that billionaires' taxes are unfairly low by approximately 0.22 standard deviations. However, rather than increasing demand for redistribution, the treatment reduces the desired income tax rate by about 3.9 pp, lowers the desired corporate tax rate by 3.9 pp, and decreases overall policy support by 0.11 standard deviations. We propose tax-hike aversion as one possible explanation for this unexpected finding. In a nutshell, the perceived current tax rate serves as a reference point: the farther a proposed rate lies above that reference point, the more reluctant individuals are to support it. Individuals may want billionaires to pay higher taxes but nevertheless be reluctant to endorse an increase that feels too large relative to the current rate. As a result, they may settle for a rate between the perceived status quo and their ideal rate. We show that, consistent with the model of tax-hike aversion, the negative effect on desired tax rates and positive effect on perceived unfairness are concentrated among individuals who initially overestimated billionaires' tax rates. Intuitively, individuals who initially overestimated billionaires' tax rates learn that actual rates are low and therefore come to view them as more unfair. At the same time, however, they demand more moderate tax increases in order to avoid endorsing a large tax hike. We also consider alternative explanations involving tax avoidance, enforceability, taxpayer mobility, and backfiring in response to distrusted information, and present suggestive evidence against these alternative channels.

In a sub-treatment, individuals receive additional information on the accounting strategies used by billionaires to reduce their tax burden. Although at baseline, most people already believed that billionaires abuse the tax system, this treatment increases that belief even further. This treatment also has a strong negative effect on non-tax attitudes toward billionaires, such as their perceived trustworthiness. The treatment does not significantly affect subjects' preferences regarding top income or corporate tax rates, but it does increase

overall policy support by 0.08 standard deviations. Our preferred interpretation is that individuals see little value in raising tax rates—since they can be circumvented anyway—but become more receptive to addressing the root of the problem through broader tax reform.

We present several robustness checks. First, we show that most of the effects persisted in the follow-up survey conducted a month later, although at about half their original size. Second, we show that none of the results are driven by any single billionaire. And finally, we show that the results are robust across alternative specifications.

We conclude the paper with a discussion of the limitations of our findings with respect to external validity. Our findings should be interpreted as evidence about the public images of five highly visible U.S. billionaires in a U.S. setting, rather than as evidence about all billionaires, wealthy individuals more broadly, or other countries and tax systems. Moreover, our treatments capture only selected dimensions of billionaires' public image; information about other dimensions, such as philanthropy, product innovation, or frugality, may generate different responses.

This paper relates and contributes to several strands of literature. Most directly, this paper relates to studies on the demand for taxation at the top of the income distribution. In particular, some studies explore the preferences for redistribution from business oligarchs (Di Tella et al., 2021), the top-1% (Hope et al., 2023), and CEOs (Kiatpongsan and Norton, 2014).[2] We contribute to this literature by examining public demand for taxing billionaires and the companies they have founded. These topics are important not only because of the substantial potential for tax revenue but also due to the growing economic and political influence that billionaires and their companies exert on the world's future. Our contribution lies in leveraging existing insights and applying them to the context of billionaire taxation—an important yet surprisingly underexplored area of research. Most closely related, the global survey conducted by Cappelen et al. (2025) documents broad support across 40 countries for a globally coordinated wealth tax on billionaires, examines how this support varies across countries and socioeconomic and political groups, compares it with support for broad domestic redistribution, and relates it to government positions. Our study complements theirs in three main ways. First, while Cappelen et al. (2025) focus on cross-country evidence concerning support for a specific globally coordinated policy—neither cross-country variation nor international coordination is part of our analysis—we study other potential drivers of the demand for billionaire taxation, such as perceived lavishness and the perceived role of merit. Second, we provide causal evidence through randomized information treatments tied to the public im-

[2]Beyond economics, there is related work in sociology and political science using mixed methods, such as open ended surveys and hypothetical conjoint experiments, to investigate how perceptions of the attributes and personality traits of the rich shape demand for taxation (Ragusa, 2015; Kane and Newman, 2023; Trump, 2025).

ages of specific, well-known billionaires, whereas their evidence on support for the billionaire tax is primarily descriptive and associational. Third, we examine a broader range of outcomes than support for a specific policy proposal, including preferred personal and corporate tax rates, perceived fairness, and behavioral measures. Our findings show that some conventional mechanisms from the broader literature on redistribution preferences—such as the distinction between merit and luck—do not explain support for billionaire taxes. Meanwhile, perceived lavishness—a channel that has received limited attention in prior research—proves to be the most effective. More generally, our work also connects to the broader literature on demand for redistribution and the role of misperceptions (e.g., Cruces et al., 2013; Kuziemko et al., 2015; Hauser and Norton, 2017).

2 Experimental Design and Implementation

2.1 Overview of the Research Design

The samples of the baseline and follow-up survey instruments are attached as Appendix E and Appendix F, respectively. Moreover, the structure of the surveys is summarized as a flow chart in Figure D.1.

In the baseline survey, each participant is randomly assigned to one of five billionaires. The survey begins with a brief introduction to the billionaire chosen for the respondent, as well as one of the companies they are best known for: Elon Musk (founder of Tesla), Jeff Bezos (Amazon), Bill Gates (Microsoft), Mark Zuckerberg (Meta), and Michael Bloomberg (Bloomberg L.P.). The five billionaires used in our experiment were selected from the top ten richest individuals in the world according to Forbes World's Billionaires List of 2023.

We deliberately chose to focus on real, well-known billionaires rather than hypothetical or generic profiles. First, these prominent individuals are particularly interesting to study because they lead highly public lives and likely play a disproportionate role in shaping public perceptions of the ultra-wealthy. They are frequently featured in the media for both their business achievements and personal activities. For instance, Elon Musk and Jeff Bezos have been named Time magazine's Person of the Year, while Michael Bloomberg served as mayor of New York City and even ran for president. Indeed, the vast majority of our survey participants reported being familiar with these billionaires.[3] The companies associated with them also rank among the most recognizable and valuable in the world. Focusing on specific billionaires, rather than a generic one, offers a second key advantage: it allows us to tailor the information treatment to real individuals, making the content more tangible and credible

[3]For more details, see Appendix D.1.

for participants. For example, a generic message stating that billionaires "got lucky" may be far less persuasive than a narrative about how specific billionaires built their fortunes, supported by journalistic accounts.

The subjects are randomized into different treatment arms. Within each treatment arm, we conduct an information-provision experiment: each subject is randomly assigned to receive a piece of information (treatment group) or to not receive any information (control group). Each information treatment provides a narrative that can contain not only numerical information (e.g., the hourly earnings of the billionaire) but also other elements, such as a picture of the billionaire's house or a screenshot of a newspaper headline. In all of the treatments, we explicitly provide a source for the information, typically a newspaper article, with a link to it. To assess whether an information treatment had any effect on perceptions, we ask a question related to the information contained in the message, both before the information-provision stage (prior belief) and after the information-provision stage (posterior belief). For example, the luxury message mentions, among other things, the estimated price of the billionaire's home. In the prior and posterior beliefs, we ask the subjects to guess the price of the billionaire's home. This will allow us to document some of the initial misperceptions, as well as whether they updated their perceptions in the expected direction. After the elicitation of posterior beliefs, the next block of questions, which is identical across all subjects, includes the outcome variables. Finally, at the very end of the survey, we collect some standard background information about the subject, such as demographics.

We took several steps to prevent respondents from making unintentional inferences due to being provided with information. For instance, respondents might assume they were chosen to receive information because their prior belief was inaccurate. To mitigate this concern, we made the randomization process explicit: we first informed respondents that some participants would be randomly selected to receive information and that they would find out on the next screen if they were chosen. Another concern is that when subjects are asked about their posterior beliefs, they might think the repeated question indicates their initial response was incorrect. To address this, we clearly informed respondents that all survey participants are asked the same question twice, regardless of their initial guesses or whether they received information.

2.2 Treatment Arms

We designed four different treatment arms, summarized below. Some treatments were inspired by previous research on the demand for redistribution and the arguments made by academics, journalists, and the general public advocating for taxing the ultra-wealthy. Figure 1 provides a sample screenshot of each information treatment. Although panels (e) and

(f) are common to all billionaires, panels (a) through (d) are specific to the billionaire selected for the respondent. In Figure 1, we use Bill Gates as an example. The corresponding screenshots for each of the other four billionaires are reported in Figures D.2 to D.5.

Luxury Treatment Arm: Motivated by the observation that some billionaires deliberately cultivate a frugal image (Del Valle, 2018), we investigate public perceptions of their consumption habits and examine how people respond to evidence of lavish spending. Specifically, we provide a picture of one of the luxurious homes purchased by a billionaire, along with information about the home's price and features. Panel (a) of Figure 1 shows a sample screenshot for Bill Gates, noting that he owns a $130 million mansion in Washington, which boasts extensive amenities such as luxurious pools, a movie theater, and a reception area that can accommodate 200 guests. In this and other treatment arms, we used neutral language in the information treatments by providing factual information (e.g., a picture, a figure) and refraining from endorsing any policies explicitly.[4]

Non-Merit Treatment Arm: According to a large body of work on social preferences (Cappelen et al., 2007; Durante et al., 2014; Almås et al., 2020; Cohn et al., 2023), individuals are more inclined to reduce inequalities attributed to luck (e.g., a coin toss) rather than those attributed to merit (e.g., performance in a task). Likewise, laboratory evidence indicates that individuals are more likely to redistribute when they believe allocations resulted from cheating (Bortolotti et al., 2023). In this arm, we inform subjects that factors beyond the billionaires' own merit played a significant role in their fortune. For example, panel (b) of Figure 1 shows the message for Bill Gates, which references a newspaper article claiming he stole the idea for Microsoft. The message thus suggests that his success was not solely attributable to merit—that is, to honest hard work. For one of the billionaires, Michael Bloomberg, we were not able to find an article from a reputable source that would be a good fit for this treatment. For that reason, in this treatment arm, subjects are randomized to one of the other four billionaires.

Earnings Treatment Arm: We teach subjects about how much billionaires actually earn. When arguing for higher taxation for billionaires, supporters of these policies often mention the unthinkable levels of wealth amassed by billionaires (e.g., Zucman, 2024). According to a body of work on social preferences, some individuals are averse to inequality (Fehr et al., 2026). Additionally, research shows that individuals have significant misperceptions about the income distribution (Cruces et al., 2013; Kuziemko et al., 2015). In this treatment arm, we provide a side-by-side comparison between the earnings of billionaires and

[4]Despite our efforts, the mere fact of studying the taxation of the ultra-wealthy may be perceived as partisan. At the end of the follow-up survey, we asked respondents whether they thought the survey was biased: 71.4% said the survey was unbiased, 25.8% said it had a left-wing bias, and 2.8% said it had a right-wing bias.

the earnings of the average American.[5] One potential concern we had with this treatment arm was that billionaires' earnings may be so astronomical that individuals are unable to grasp their true magnitude. For example, individuals rarely have to deal with billions of dollars in their day to day lives. To address this concern, subjects are randomized to one of two conditions: the main condition of hourly earnings (assigned with $\frac{2}{3}$ probability) and the alternative condition of annual earnings ($\frac{1}{3}$ probability). The hourly sub-treatment is identical to the annual sub-treatment, except that it includes additional text converting the annual salaries into their hourly equivalents. For the sake of brevity, and to maximize statistical power, the main specification pools these two conditions—in any case, we do not find significant differences between them (results reported in Appendix D.2). Panel (c) and (d) of Figure 1 show screenshots of the hourly and annual treatments for Bill Gates. While the typical American earns $64,000 per year, Gates earns $3.05 billion per year. For the typical American, if you divide the annual income of $64,000 by the 8,760 ($= 365 \cdot 24$) hours in a year, it comes out to about $7 per hour. For Bill Gates, the corresponding estimate would be $348,015 per hour.

Tax Treatment Arm: There is evidence that, when forming preferences over taxation and deciding whether to comply with taxes, individuals care about fairness (Stantcheva, 2021; Nathan et al., 2023). According to some accounts, billionaires pay lower effective tax rates than the typical American (e.g., New York Times, 2019; Leiserson and Yagan, 2021; ProPublica, 2023; EU Tax Observatory, 2024; Zucman, 2024). If the public learns about this fact, they may conclude that billionaires are not paying their fair share. In the tax rate treatment, we provide a side-by-side comparison between the tax rate paid by the typical American and the tax rate paid by the 25 richest billionaires. A screenshot of the treatment is shown in panel (e) of Figure 1: according to ProPublica (2023), billionaires paid an average tax rate of 16%; in comparison, the average American pays a tax rate of 21%.[6] It is possible that individuals do not mind that billionaires face low tax rates if they play by the rules, but they may react differently if they perceive that billionaires are cheating their way into low taxes. For example, evidence from laboratory experiments shows that individuals are more willing to redistribute resources from those perceived to have cheated (Di Tella et al., 2015; Bortolotti et al., 2023). To test this additional hypothesis, we cross-randomized an additional message. A screenshot of the treatment is shown in panel (f) of Figure 1. In addition to the tax rate information, some individuals receive an additional screen with information on two common accounting strategies used by billionaires to reduce their tax burden: receiving

[5] In the question about prior beliefs, we informed subjects about the earnings of the typical American, so this information was available both in the treatment and control groups.

[6] In the question about prior beliefs, we informed subjects about the tax rate paid by the typical American, so this information was available both in the treatment and control groups.

compensation in stocks and borrowing against them to avoid taxes until sold, and setting up smaller companies in tax havens to transfer profits and minimize tax liabilities.

In each of the treatment arms, we asked for prior and posterior beliefs related to the information provided in that arm. All of these questions are listed in Table D.1 and summarized below. In the luxury treatment arm, we asked subjects to guess the value of the billionaire's home.[7] In the non-merit treatment arm, we elicited prior and posterior beliefs by asking the extent to which the respondent agrees with the statement that the billionaire "earned his wealth through honest and hard work." In the earnings treatment arm, we asked the respondent to guess the earnings of the billionaire (hourly or annual, depending on the subtreatment assigned to the respondent). In the tax treatment arm, we included two questions for prior and posterior beliefs. The first question, designed with the tax rate information in mind, asked the subject to guess the effective tax rate paid by the billionaire to which they were assigned. The second question, designed with the information about the tax loophole in mind, asked subjects to determine to what extent they agree with the statement that "billionaires abuse loopholes in the tax code to avoid paying taxes."

2.3 Outcomes of Interest

At the end of the survey, we included a series of questions that make up the outcome variables and are listed in Table D.2. The outcome variables can be categorized into three broad groups, as summarized below.

Taxation Preferences: The primary outcome of interest is the top income tax rate for billionaires. This outcome is inspired by real-world policy proposals to establish a top income tax rate for billionaires, such as those proposed by President Biden, Vice-President Harris and Congresswoman Ocasio-Cortez.[8] We ask respondents to imagine that the U.S. government is planning to introduce a new personal income tax rate specifically for individuals earning more than $10 million annually, and ask them to use a slider to select a marginal tax rate between 0% and 100%.[9] The next question is intended to examine whether perceptions about billionaires influence the demand for taxation of the companies they founded. Respondents

[7] Given that the amounts individuals could enter are so large, there is a concern for large outliers due to typos. For that reason, in this and other numerical questions, we included a numerical validation—for more details, see Appendix B.2.

[8] The government has other mechanisms to increase taxes on the ultra-wealthy, such as raising the top capital gains tax. Given that billionaires can avoid income taxes, raising the top income tax rate may not be the most effective policy (see e.g., The Economist, 2024). However, we considered this an appropriate outcome because, while not all subjects pay capital gains taxes, a strong majority are subject to income taxes and thus are probably familiar with them.

[9] To avoid influencing their choices, this and other sliders do not have a default position. The subject needs to click somewhere on the horizontal line for the slider bar to appear.

are asked to imagine that the U.S. government is planning to introduce a corporate tax for companies that make profits exceeding $10 million—we explicitly mention that this would include large businesses such as the one owned by their assigned billionaire. Subjects can use a slider to select a corporate tax rate between 0% and 100%. Using survey questions to elicit desired taxes or tax rates is a well-established method in the literature for measuring attitudes toward redistribution and tax policy (see e.g., Kuziemko et al., 2015; Fisman et al., 2020; de Bresser and Knoef, 2022; Alesina et al., 2023). Recognizing that individuals may want to tax billionaires for reasons of efficiency or fairness, we also include a question probing the fairness aspect directly. This question asks whether, from a fairness perspective, the taxes paid by the billionaire are too high or too low on a 7-point scale.

Policy Support: We construct an index of policy support based on six different metrics and standardize it so that the control group has a mean of 0 and a standard deviation of 1.[10] We describe four real policy proposals and for each of them, we elicit the respondent's support on a 7-point scale from strongly oppose to strongly support. The four proposals are: (i) a proposal that would require the wealthiest American households to pay a minimum of 20% of their total income in taxes; (ii) the introduction of an international corporate tax at an annual rate of 0.2% of the company's market value; (iii) a proposal to introduce a wealth tax for individuals with wealth over $50 million;[11] (iv) a newly introduced law that allocated $80 billion in funding to the IRS to strengthen tax enforcement.[12] To incentivize subjects to answer truthfully, we informed them that we would share the anonymous survey results with politicians and relevant organizations. Previous studies found this approach to increase subjects' perceptions of how consequential their responses are (e.g., Elías et al., 2019). Additionally, we elicited two revealed-preference measures of policy support. First, we allow individuals to allocate a $300 budget for donations to two organizations: World Relief, a Christian non-governmental organization that provides humanitarian aid and development assistance to vulnerable communities around the world; and Americans for Tax Fairness, which advocates for big corporations and the wealthy to pay their fair share in taxes. To give subjects incentives to be truthful, we tell them we will split the donations according to the choices of one randomly chosen respondent.[13] Second, respondents are given the opportunity to sign a petition organized by Oxfam to increase taxes on the ultra-rich. We

[10] We standardize all the individual items of the index using the mean and standard deviation of the control group. To create the index, we take the average of the standardized items and then standardize this average again using the mean and standard deviation in the control group.

[11] While this tax was proposed for the state of California, we asked subjects whether they would support a similar policy in their own state of residence.

[12] We randomized the order in which subjects saw the proposals.

[13] We chose the lesser-known NGO World Relief with the aim of balancing choices and avoiding ceiling effects, thereby allowing us to detect treatment effects in both directions. This design choice proved effective: in the control group, the median donation split was close to fifty-fifty.

asked respondents whether they want to sign the petition and, if they say yes, we show them a screen with a link to sign it.[14]

Non-Tax Attitudes: Providing information about basic aspects such as how much billionaires pay in taxes or how truly wealthy they are may also change attitudes toward them that are unrelated to taxation. While these attitudes are not an outcome of interest in itself, they can shed light on the causal mechanisms behind the effects on preferences for taxation. To achieve that goal, we construct a non-tax attitudes index based on five different metrics. Like in the policy index, we construct a standardized index with a mean of 0 and a standard deviation of 1. We ask respondents whether they believe the billionaire deserves his wealth, whether they find the billionaire trustworthy, and if they admire and respect him. Additionally, we measure respondents' feelings towards the company founded by the billionaire. First, we use a subjective question to assess whether the individual has a positive attitude towards the company. Second, we use a revealed-preference method. Subjects are shown a picture of a backpack with the logo of the company. Using a multiple price list method, we elicit the willingness to pay for the backpack. Intuitively, subjects with a negative image of the company should not want to walk around with the company's logo on their back and, therefore, should be willing to pay less for the backpack.

2.4 Background Characteristics

The final block of questions collected background information about the respondents, which can be used for descriptive analysis, as control variables, and for heterogeneity analysis. We included a standard set of questions regarding the gender, age, ethnicity, income, and education of the subjects. Additionally, we elicited the subjects' partisan identity. We also assessed subjects' general attitudes towards redistribution on an 11-point scale ranging from no redistribution to complete redistribution. Furthermore, we measured trust in the federal government on a 4-point scale from low trust to high trust.

2.5 Follow-up Survey

To assess the persistence of treatment effects, we invited all subjects to complete a follow-up survey approximately one month after the baseline. This follow-up survey measured the same posterior beliefs and outcome variables collected in the baseline survey, with just two minor

[14] In addition, we collect some data to validate the petition outcome. The results are presented in Appendix D.4.

exceptions.[15] In an effort to mitigate experimenter demand effects, we took several measures to obfuscate the connection between the baseline and follow-up surveys (Haaland et al., 2023). First, we used different Prolific accounts for the surveys: subjects were invited to the baseline survey from a UC-Berkeley account and to the follow-up survey from a University of Zurich account. Second, we changed the layout of the follow-up survey, such as using different fonts and colors, and replaced the UC-Berkeley logo with the University of Zurich logo. Additionally, to further obfuscate the connection to the baseline survey, the follow-up survey began with a series of filler questions about the subjects' use of and attitudes towards generative artificial intelligence.

2.6 Implementation of the Experiment

We conducted the pre-registered survey with a sample of 9,013 Americans recruited via Prolific.[16] We advertised the study on Prolific as "Scientific study" with an estimated duration of 10 minutes and a participation reward of $2. We collected responses in two waves, in January and March 2024.[17] About 82% of these participants also completed the follow-up survey, which had an estimated duration of 5 minutes and offered a participation reward of $1.50. The median time elapsed between the baseline and the follow-up survey was 24 days.[18] We limited our participant pool to U.S. residents, adhering to best practices for recruiting and ensuring high-quality responses. In both the baseline and follow-up survey, we included Captcha verification. Furthermore, we included different attention checks in both surveys, which 99% of subjects passed—subjects who did not pass the attention checks were excluded from the analysis. The median completion times were 9 minutes for the baseline survey and 6 minutes for the follow-up survey. At the end of the baseline survey, we asked respondents about the difficulty of the survey, and 89% indicated that our survey was "easy to understand."

2.7 Descriptive Statistics and Balance Checks

Table D.3 presents descriptive statistics. Approximately 49% of the subjects are female, 47% are 35 years old or younger, 65% have an annual household income above $50,000, 69% have

[15] In the follow-up survey, we did not ask subjects if they wanted to sign the petition again, as one cannot sign a petition twice (we did ask subjects if they had heard of the petition before, for a validation check). Second, to keep the follow-up survey short, we did not elicit the willingness to pay for the backpack.

[16] The experiment was pre-registered at the AEA RCT Registry (AEARCTR-0011845).

[17] The first wave was open from January 27 to 31 of 2024, with a total of 6,016 responses. The second wave was open from March 28 to 30 of 2024, with 2,997 responses.

[18] For each wave, we published the follow-up survey three weeks after completing the baseline survey, keeping it online for 35 days in wave 1 and 21 days in wave 2, respectively. In both waves, 90% of the follow-up data was collected within the first week.

a college degree, 63% self-identify as White, 15% as Black, 10% as Asian, and 8% as Hispanic. Our sample is not perfectly representative of the general U.S. population, but it is not very dissimilar either—for more details, see Appendix D.3. As is common in online samples, the subject pool skews toward younger, more educated, and left-leaning individuals. The most notable imbalance concerns political affiliation: approximately 50% of respondents identify as Democrats, compared with 35% in the representative benchmark reported in Table D.4. The average treatment effects apply directly to our subject pool and could differ from those for the broader U.S. population if political affiliation moderates responses to the treatments. We return to this issue in Section 3.6, where we examine heterogeneity in treatment effects by political affiliation.[19]

Table D.3 shows that, consistent with successful randomization, the pre-treatment characteristics are balanced across treatment groups.[20] Table D.3 also shows the participation rates in the follow-up survey. On average, 82% of the subjects who participated in the baseline survey also participated in the follow-up survey. Participation rates are similar across treatment conditions in the luxury, non-merit, and earnings arms. In the tax arm, however, there is a statistically significant difference across conditions: participation rates are similar among respondents who received only the tax rate treatment (83.8%) and those in the control group (82.2%), but lower among respondents who received the tax loophole treatment (78.1%). Thus, we provide robustness checks to address concerns that this differential attrition could bias the follow-up estimates. Additionally, we observe virtually no attrition within the baseline survey.[21]

2.8 Supplemental Emotions Survey

To explore subjects' emotional reactions to the information treatments, we conducted a supplemental survey referred to as the "emotions survey."[22] We present the same information about billionaires as in the baseline survey, but then elicit the respondents' emotional reaction to the treatments using a combination of multiple-choice and also open-ended questions (e.g., Haaland et al., 2024). Subjects were randomized into one of the information treatments,

[19] Although political affiliation strongly predicts baseline support for redistribution, this does not necessarily imply heterogeneity in treatment effects, as respondents with different initial preferences may respond similarly to the same information.

[20] The differences between treatment and control groups are always small in magnitude. Four of the differences are statistically significant at the 10% level. However, those differences are probably spurious: given that 60 tests are reported in this table, we expect that about 6 of them should be statistically significant at the 10% level just by chance.

[21] 98.5% of the subjects who started the survey completed it. And conditional on reaching the information-provision stage, 99% of participants completed the survey.

[22] We provide more details on the emotions survey in Appendix C.2.

after which they reported their thoughts and feelings in an open-ended response and selected experienced emotions from a predetermined list. We then analyzed the open-ended responses using a large language model (LLM) to identify the two most common response categories for each treatment. Subsequently, two research assistants independently assigned each response to one of the two main categories or to a residual category.

3 Results

3.1 Baseline Attitudes

We start by describing the baseline beliefs, preferences, and attitudes in the control group.

Figure 2 presents histograms with a selection of key outcomes of interest—the rest of the outcomes are presented in Figures D.7 and D.8. This evidence shows that while there is considerable variation in support for taxing billionaires, preferences are skewed in favor of higher taxation. For example, panel (a) of Figure 2 shows that the average participant prefers a top income tax rate of 42.5%; for reference, this rate is not substantially higher than the current top income tax rate of 37%. Moreover, there is substantial heterogeneity in the preferred tax rates, ranging from a 10th percentile of 16% to a 90th percentile of 75%. In turn, panel (b) of Figure 2 shows the distribution of the preferred top corporate tax rates. The results mimic those of the income tax rate: on average, individuals prefer a 39.8% corporate rate, but there is large variation between individuals.[23] The average preferred rate of 39.8% is substantially higher than the current flat rate of 21%, and even higher than the maximum rate of 35% that was effective before the Tax Cuts and Jobs Act of 2017. Support for taxation is even stronger when participants are asked explicitly about fairness: they are much more likely to say that the taxes billionaires pay are too low than to say that they are too high, as shown in panel (c) of Figure 2.

Furthermore, a majority of subjects support the policy proposals aimed at increasing the taxation of the ultra-wealthy. For example, panel (d) of Figure 2 shows that 78.5% support the billionaire minimum income tax proposal. The other three policy proposals received similar support.[24] The support for policy change is also reflected in the revealed-preferences measures. For example, panel (e) of Figure 2 shows that when asked to divide the $300 donation between two charities, a strong majority (84.7%) allocated at least some amount

[23] Appendix Figure D.9 presents the distributions of preferred income and corporate tax rates in greater detail using one-percentage-point bins. Both distributions exhibit bunching at multiples of five, indicating that respondents often report rounded values that are best interpreted as coarse measures of their underlying preferences. Such bunching is common in survey responses elicited in percentage points, including reported probabilities (e.g., Manski and Molinari, 2010).

[24] For more details, see Figure D.7.

to Americans for Tax Fairness. Furthermore, most subjects (59.7%) are willing to sign a petition to increase taxes on the ultra-rich.

The substantial cross-sectional heterogeneity we document in respondents' preferred tax rates for billionaires (and for corporations) could, in principle, be driven partly—or even largely—by measurement error. To test this hypothesis directly, we use a standard tool from the survey literature: the test-retest correlation. Consider individuals in the control group. These respondents are asked the same question about their preferred tax rates twice—once in the baseline survey and again in the follow-up survey. Intuitively, we would expect them to provide the same answer both times, since they received no new information as part of the experiment and therefore have no reason to change their views. If, instead, their responses were purely noise, we would expect the two answers to be uncorrelated. The strength of the correlation between the two responses therefore indicates whether we are closer to one extreme (no measurement error) or the other (pure measurement error).[25]

We find that our main outcome variables are highly correlated across survey waves, with correlation coefficients ranging from 0.683 to 0.760.[26] While a high test-retest correlation does not in itself validate a question, the fact that responses remain highly stable over time strongly suggests that the measure is capturing something meaningful rather than mere noise. For comparison, these test-retest correlations are nearly as high as those for hard outcomes such as education and income, which can reach 0.90 (Krueger and Schkade, 2008). Moreover, the four outcomes in this study exhibit higher test-retest correlations than those reported for life satisfaction—about 0.59 (Krueger and Schkade, 2008)—and for other economic preference measures (Dohmen and Jagelka, 2024).

In terms of non-tax attitudes towards the billionaires and the companies they founded, there is considerable variation across subjects, but the non-tax attitudes are slightly skewed in the positive direction. For example, panel (f) of Figure 2 shows the responses to the question on whether billionaires deserve their wealth. On a scale from "not deserving" (0) to "deserving" (6), the average response is 3.3, which is slightly closer to the positive end of the spectrum. The respect for billionaires and the attitude towards their companies are also skewed positively; however, billionaires score consistently low in measures of trustworthiness.[27]

We also find that the subjects had significant misperceptions about the billionaires. Three of the information treatments include some numerical facts for which we elicited prior beliefs.

[25] The test-retest framework applies most cleanly with an immediate follow-up. In our setting, the follow-up occurs about one month later, during which respondents may legitimately revise their preferences. If anything, this reduces the correlation, so our estimate should be interpreted as a conservative lower bound on the true correlation.

[26] For more details, see Figure D.22.

[27] For more details, see Figure D.8.

By comparing these prior beliefs to the facts provided in the messages, we find significant misperceptions.[28] In the luxury treatment arm, we asked the subject to guess the value of one of the billionaire's homes. Only a small share (3.3%) of the participants could correctly guess the value of the home (i.e., within ± 5% of the truth), and most subjects were off by a wide margin.[29] Furthermore, subjects were more likely to under-estimate than to over-estimate this value.[30] Likewise, only a small minority of individuals (0.8%) could correctly guess the earnings of the billionaire (i.e., within ± 5% of the truth), again with more individuals under-estimating than over-estimating.[31] Lastly, we asked about the effective tax rate faced by the billionaire. The true tax rates that each billionaire pays are not publicly available, so we cannot compare the guesses to the true rate. However, we can compare the guesses to the average tax rate for the top 25 richest billionaires estimated by ProPublica (2023).[32] Only a small share (1.2%) of guesses came close (within ± 2.5 pp) to ProPublica's estimate. Moreover, subjects are more likely to guess above than below ProPublica's estimate.

Finally, we examine how prior beliefs relate to our main outcome, the preferred tax rate (see Appendix D.6 for details). In all treatment arms except the annual earnings treatment, prior beliefs are significantly correlated with the preferred tax rate, with signs in the expected directions. For example, believing that a billionaire's home is more expensive is associated with a higher desired tax rate, whereas believing that a billionaire's success is due to honest and hard work is associated with a lower desired tax rate. However, these correlations may be driven by omitted variable bias. The survey experiment therefore provides causal evidence on whether these beliefs affect individuals' demand for taxation.

3.2 Luxury Treatment Arm

We begin by documenting the effects of the luxury treatment on beliefs. Figure 3 shows a histogram of the distribution of posterior beliefs. Each panel corresponds to a different treatment arm. In each panel, the red bins represent individuals in the treatment group (who received the information), while the gray bins correspond to individuals in the control group

[28] In all the results presented in this paper, we mitigate sensitivity to outliers by winsorizing prior and posterior beliefs about the value of the billionaire's home, hourly and annual earnings, and the total tax rate. The threshold used for winsorization is based on the 95th percentile of absolute prior misperceptions.

[29] To make it directly comparable to the rest of the statistics reported in this subsection, this and the next statistics correspond to individuals in the control groups. However, since this is a pre-treatment outcome, the results are almost identical for individuals in the treatment group.

[30] For more details, the full distributions of prior beliefs are reported in Figure D.10.

[31] The result on under-estimation, however, is clear when the earnings were elicited hourly but not clear when they were elicited annually.

[32] This comparison has to be taken with a grain of salt, because the estimate is subject to error. Moreover, there may be significant differences between how much each of the billionaires pay in tax rate and the average for the top-25.

(who received no information), except in panel (f), which compares the two sub-treatments. In panel (a), for the luxury treatment arm, the x-axis shows the difference between the individual's guess for the value of the billionaire's home and the true value (i.e., according to the treatment). For example, a value of \$0 means that the guess is accurate, while a value of -\$1 million means that the individual under-estimated the home value by \$1 million. In the control group, only a small minority of subjects have accurate posterior beliefs. In contrast, in the treatment group, a large majority of subjects have accurate posterior beliefs. The fact that subjects in the treatment group have more accurate beliefs suggests that subjects paid attention to the information and found it reliable.[33]

Evidence from the supplemental emotions survey indicates that most participants engaged emotionally with the luxury message. In the multiple-choice responses reported in Figure 4, only 8.3% selected no emotion or indifference, while negative emotions were substantially more common than positive ones (66.1% versus 25.6%). The open-ended responses summarized in Table D.5 show a similar pattern: 46.8% expressed frustration or anger about extreme wealth disparities, often describing the luxury as excessive or wasteful. Given these predominantly negative reactions, one might expect the luxury message, if anything, to increase support for taxation.

Appendix D.7 provides a more detailed analysis of the effects of information. We find that individuals updated their posterior beliefs in the direction of the information provided. Individuals who initially under-estimated the value tend to update their posterior beliefs upward, while those who initially over-estimated the value tend to update their posterior beliefs downward. Despite some participants updating upward and others downward, on average, the luxury treatment leads individuals to perceive that billionaires live in more expensive homes. Panel (a) of Figure 3 shows that the treatment caused an increase of \$4.01 million (p-value=0.189) in the average belief about the billionaire's home value. For numeric elicitations like this one, a common concern is that outliers, often due to typos or misunderstanding of the question, may disproportionately influence the difference in means. For that reason, Figure 3 also reports the difference in medians. Panel (a) shows that the treatment caused an increase in the median belief of \$9.00 million (p-value<0.001).

Next, we measure the average effects of the luxury treatment on the demand for taxation. To estimate these effects, we used a simple linear regression model. Let Y_i^{post} be the outcome of interest. For example, the main outcome is the respondent's preferred top income tax rate. Let T_i be an indicator variable that takes the value 1 if the respondent was randomly

[33] As a falsification test, we also compare the distribution of prior beliefs between the treatment and control groups. Since prior beliefs were elicited before the information-provision stage, the treatment should not have any effect on them. This is confirmed by our findings, as discussed in Appendix D.5.

assigned to receive information.[34] The regression of interest is as follows:

$$Y_i^{post} = \nu_0 + \nu_T \cdot T_i + X_i^{pre}\nu_X + \varepsilon_i \tag{1}$$

ν_T is the main coefficient of interest, corresponding to the Average Treatment Effect (ATE) of the information.[35] X_i^{pre} corresponds to the vector of control variables from the module on background characteristics: gender, age, income, education, ethnicity, a dummy for Democrat, a dummy for Republican, the general attitude towards income redistribution, and the trust in the federal government.[36] Since the treatment was randomized, the control variables are not needed for causal identification. However, they can help reduce the variance of the error term and thus improve statistical precision (McKenzie, 2012).[37] In any case, as discussed below, the results are similar if we include a more limited set of controls or if we do not include any controls.

The estimated ATEs are reported in Table 1. Each row corresponds to a different treatment, and each column corresponds to a different outcome variable: column (1) corresponds to the main outcome, the preferred top income tax rate; column (2) corresponds to the preferred corporate tax rate; column (3) corresponds to perceived tax fairness; column (4) corresponds to the policy index; and column (5) corresponds to the non-tax attitudes index. Given the large number of combinations between treatments and outcomes, it is important to account for multiple hypothesis testing. For each coefficient reported in Table 1, we also provide the corresponding sharpened q-value in brackets (Benjamini et al., 2006; Anderson, 2008). The q-value represents the minimum false discovery rate (i.e., the expected proportion of rejected null hypotheses that are actually true) at which the null hypothesis would be rejected for that specific test, considering all tests reported in the same table.

The top row of Table 1 reports the results for the luxury treatment. This treatment has positive effects on the demand for taxation and across the different outcomes. Column (1) shows that, relative to the control group, individuals exposed to the luxury treatment prefer

[34] This specification applies to the treatment arms with a single treatment group. In the tax treatment arm, we include two treatment variables: one dummy indicating whether the subject received the information on tax rates, and another dummy indicating whether the subject received the additional information on tax loopholes.

[35] In our study, the primary goal is to estimate the average treatment effects of providing information. In other information-provision experiments, the primary objective is to identify the causal effect of beliefs themselves. In those settings, the impact of the information treatment would be interpreted as an intention-to-treat effect.

[36] We control for partisan identity following other studies (e.g., Alesina et al., 2023).

[37] Due to potential survey fatigue, we included the background questions at the end of the survey. In theory, the treatment could affect the responses to these questions because they were elicited after the information-provision stage. In practice, this is not a concern: Table D.3 shows that there were no treatment effects on any of the background characteristics.

a top income tax rate that is 2.0 pp higher (p-value=0.036). Column (2) shows that the treatment also raises the preferred corporate tax rate by 1.8 pp (p-value=0.053), indicating that the increased demand for taxation extends to companies founded by the billionaire. Column (3) shows that treatment increases the perception that billionaire taxes are unfairly low, by 0.12 points (p-value=0.048), equivalent to 0.082 standard deviations. Furthermore, column (4) shows a significant effect on policy support of 0.075 standard deviations (p-value=0.034).[38] All of these effects are statistically significant even after accounting for multiple hypothesis testing (q-values of 0.083, 0.095, 0.094 and 0.083, respectively).

We next turn to the economic magnitude of the effect. The luxury treatment increases our primary outcome by 2.0 percentage points, corresponding to roughly 10% of a standard deviation. This effect is substantively meaningful, especially in light of evidence that many information-provision experiments generate substantial belief updating but little to no change in policy preferences (Haaland et al., 2023). It is also notable given the light-touch nature of our intervention, which consists of a single page featuring an image and brief text.

To provide a quantitative benchmark, we compare our estimates to those in Kuziemko et al. (2015). In a survey experiment, they provided respondents with information about income inequality in the United States and examined the impact on attitudes toward redistributive policies. Across treatments, they report effects on the desired average income tax rate for the top 1 percent ranging from 5% to 12% of a standard deviation.[39] Our treatment effect on the desired top income tax rate, at roughly 10% of a standard deviation, lies toward the upper end of this range.

There are a few additional results reported in the Appendix. While Table 1 aggregates the policy and non-tax attitudes questions into two indices, Table D.7 presents the disaggregated results for each individual outcome. For example, Table 1 shows that the luxury treatment increases the policy index by 0.075 standard deviations. In turn, Table D.7 shows that the luxury treatment had a consistently positive effect on all items that comprise the index, with the most significant effects on the support for a wealth tax (p-value=0.086) and the donation to Americans for Tax Fairness (p-value=0.010).

Some individuals reacted to the luxury treatment by increasing their beliefs about the

[38]Column (5) shows that the effects on non-tax attitudes are close to zero (-0.022 standard deviations), precisely estimated and statistically insignificant. That is, when an individual is provided with information about a billionaire's lavish home, it does not change how they feel about the billionaire, such as whether they respect him. Evidence from the emotions survey, however, indicates that the treatment did elicit frustration about excessive spending and inequality, which may not be fully captured by the non-tax attitudes index.

[39]We use the data from the replication package of Kuziemko et al. (2015) to compute the reported effect sizes in standard deviation units. The 5% effect size corresponds to the treatment effect on the top 1 percent tax rate reported in Column (1) of Table 5. The 12% effect size corresponds to the treatment effect of the estate-tax-only treatment on the top 1 percent tax rate reported in Column (6) of Appendix Table 13 in the Online Appendix.

value of the billionaire's home, while other individuals updated their beliefs downward. The treatment did not only include the value of the home, but also an image, a description, and a narrative that connected all the pieces of information. Thus, the effect of this treatment may be driven by the belief about the home's value or by the other pieces of information. If the effects of the treatment operated solely through the belief about the home value, we would expect asymmetric effects: individuals who updated their beliefs upwards should react in the opposite direction than those who updated their beliefs downward. Table D.9 replicates Table 1, except that it breaks down the sample by subjects who started under-estimating (and thus will update their beliefs upwards) versus those who over-estimated (and thus will update beliefs downwards). There are no significant differences in treatment effects between these two groups. This evidence suggests that the effect of the treatment was not primarily driven by quantitative information about the value of the home. Instead, our preferred interpretation is that the treatment effects are due to the qualitative information, such as the impressive picture of the billionaire's residence and the description of the luxurious amenities. For instance, individuals may not react strongly to the numerical information because home values are so far beyond their experience that they struggle to differentiate between a $10 million home and a $100 million home. Our evidence is consistent with other survey experiments showing that, when it comes to policy preferences, individuals can be more responsive to qualitative than to quantitative information (e.g., Rasooly, 2024); and qualitative anecdotes and narratives can influence redistributive preferences more strongly than factual information (see e.g., Kuziemko et al., 2015; Alesina et al., 2023). For example, Kuziemko et al. (2015) show that combining emotional content, including pictures, with neutral information about the estate tax has a substantially larger effect on support for increasing the estate tax than providing neutral information alone.[40] However, the luxury treatment is a bundle: it provides an image, describes expensive amenities, reports the value of the home, and highlights the billionaire's consumption choices. Our design therefore cannot distinguish the relative contribution of each component. For example, respondents may react to the vivid presentation of extreme wealth, to moral judgments about conspicuous consumption, or to a combination of both. With this caveat, the findings from the luxury treatment and the supplemental emotions survey provide suggestive evidence that individuals' demand for taxation may depend on how the rich are perceived to use their resources. This interpretation is consistent with evidence from Hansen (2023) and Trump (2025), which

[40] An avenue for future research would be to examine this channel more precisely by cross-randomizing information on both the level and the use of wealth, thereby isolating the effect of conspicuous consumption. A complementary design could compare the luxury treatment with information about earnings or wealth presented in an equally vivid visual format, thereby helping distinguish the effect of visual salience from the content of the information.

shows that support for taxation can vary depending on, for example, whether the rich are described as charitable.

3.3 Non-Merit Treatment Arm

The non-merit treatment reduces the belief that honest and hard work played a significant role in the billionaire's wealth. Panel (b) of Figure 3 shows the distribution of posterior beliefs that honest and hard work played a significant role in the success of the billionaire, on a scale from 0 to 6. The non-merit treatment shifts the distribution to the left. More precisely, the information causes a reduction in the average belief of 0.9 points (p-value<0.001), or equivalent to 0.46 standard deviations. This strong effect suggests that subjects paid attention to the information and found it reliable. If the non-merit treatment persuaded individuals that the billionaire's success goes beyond his honest and hard work, it may change the attitudes towards the billionaire. We can test this hypothesis by examining the effects on the non-tax attitudes index. The second row of Table 1 reports the average treatment effects of the non-merit treatment. Column (5) shows that this treatment reduced the non-tax attitudes towards the billionaire by 0.135 standard deviations (p-value<0.001). This negative effect was prevalent across all items in the non-tax attitudes index, but was most significant for feelings of deservingness of wealth, respect, and attitude towards the company (for details, see Table D.7).

The supplemental emotions survey again indicates that, in addition to paying attention to the information, most respondents engaged with it emotionally. The non-merit message elicited more negative emotions (52.5%) than positive ones (27.2%)—see Figure 4 for details. Analysis of the open-ended responses indicates that the most common response, provided by 42.9% of participants, is to become more skeptical of claims of self-made success among billionaires (see Table D.5 for details).

Despite the strong negative effect on the perception that the billionaire achieved his success through honest and hard work, the non-merit treatment did not have any significant effects on the demand for taxation. The coefficients are mostly close to zero and always statistically insignificant: column (1) shows an effect of -0.144 pp (p-value=0.885) on the preferred income tax rate; column (2) shows an effect of 0.730 pp (p-value=0.450) on the preferred corporate tax rate; column (3) shows an effect on the belief that taxes are unfairly low of 0.089 points (p-value=0.123), equivalent to 0.060 standard deviations; and column (4) shows an effect on policy support of 0.016 standard deviations (p-value=0.649).[41] Although for specific outcomes we sometimes cannot rule out small or modest effects, taken together,

[41] We do not find any significant differences between individuals with different prior beliefs about the role of honest and hard work—results reported in Table D.9.

the coefficients suggest that the treatment was largely ineffective in increasing demand for taxation.

This result contradicts experts' predictions that the non-merit treatment would substantially increase the demand for taxation. This finding also goes against evidence from laboratory studies showing that individuals are more willing to redistribute resources when outcomes are determined by factors beyond merit, such as a coin flip. Indeed, that research served as motivation for this treatment arm and is probably what motivated experts to predict positive effects for this treatment. Our preferred interpretation is that, unlike laboratory settings, in real-world settings economic success is perceived as the result of a complex interplay of factors, including luck, skill, hard work, and opportunity. As a result, while the non-merit treatment persuades individuals that honest and hard work is not the sole driver of wealth, it may not translate into a greater demand for taxation. In fact, this interpretation is consistent with some recent laboratory evidence. For example, Yusof ⓡ Sartor (2024) shows that people tend to accept market-driven inequalities even if they are outside the individual's control. And Andre (2024) shows that individuals hold others responsible for their choices even if these choices have been shaped by unequal circumstances.

The messages also capture conceptually distinct forms of non-merit: the Bezos message emphasizes luck, whereas the others emphasize alleged appropriation or dishonesty. The design therefore provides some variation across these forms, although the billionaire and the specific narrative also vary. The estimated effect is virtually unchanged when Bezos is excluded (see Figure D.20), suggesting that the pooled result is driven neither by the luck-focused message nor by offsetting responses to the Bezos and other messages. Future research could distinguish these mechanisms more directly by randomizing respondents between different forms of non-merit information.

3.4 Earnings Treatment Arm

The earnings treatment has a significant effect on the perception of the billionaire's earnings. Panels (c) and (d) of Figure 3 show the distribution of posterior beliefs about the billionaire's hourly and annual earnings, respectively. In the control group, only a small group of subjects have accurate guesses of the billionaire's earnings. In comparison, in the treatment group, a strong majority of subjects provide accurate guesses. This strong effect of the treatment suggests that subjects paid attention to the information and found it trustworthy. Moreover, panel (c) of Figure 3 shows that individuals in the control group systematically under-estimate earnings in the main hourly condition. As a result, the treatment has a strong positive effect on both the mean and the median posterior belief, of about \$253K (p-value$<$0.001) and

$401K (p-value<0.001) respectively.[42] In the annual condition, shown in panel (d), there is a significant fraction of very large outliers that complicates the interpretation.[43]

The emotions survey suggests that the individuals engaged with the information emotionally. Although the earnings information occasionally evoked positive emotions like admiration and motivation (30.8% of subjects, from Figure 4), the majority of reactions were negative (61.3%). The analysis of the open-ended responses further supports this pattern: while some participants expressed mixed feelings, the predominant reaction was disgust and frustration over the extreme inequality (36.6% of subjects, from Table D.5).

When individuals are informed about the true earnings of billionaires, the evidence shows positive effects, although weaker, on the demand for taxation. The third row of Table 1 reports the average treatment effects of the earnings treatment. On the one hand, columns (1) and (2) show some significant positive effects on the desired income and corporate tax rates, of 1.9 pp and 1.46 pp (p-values of 0.018 and 0.064, and corresponding q-values of 0.065 and 0.106). On the other hand, we do not find statistically significant effects on the other outcomes related to demand for taxation: column (3) shows an effect on perceived tax unfairness of 0.054 points (p-value=0.267), equivalent to 0.037 standard deviations; and column (4) shows an effect on policy support of 0.015 standard deviations (p-value=0.610).[44]

Similar to the luxury treatment, the earnings treatment aims to highlight the extreme inequalities between billionaires and the average American. However, the luxury treatment has a more robust impact on preferences for taxation, significantly affecting the feeling of unfairness and the support for policy in addition to the effects on the preferred tax rates. The evidence suggests that seeing a picture of the billionaire's extravagant home can trigger a more significant reaction than merely presenting statistics.

3.5 Tax Treatment Arm

The information about the effective tax rates that billionaires pay has strong effects on beliefs and in the expected direction. Panel (e) of Figure 3 shows the distribution of posterior beliefs about the average tax rate paid by the billionaire, where individuals who received information

[42] Given that individuals are learning that the billionaires are earning more than they thought, one might expect positive effects on the attitudes towards the billionaire. Column (5) of Table 1 shows that there is a small positive effect of 0.045 standard deviations, but it is statistically insignificant (p-value=0.181).

[43] In the control group, a substantial share of respondents greatly overestimate annual earnings, likely because they confuse earnings with net worth; Appendix B.2 discusses this in more detail.

[44] Table D.9 explores the heterogeneity by prior beliefs. The results are mixed: we find some suggestive evidence that the effects on preferred tax rates are stronger for individuals with high prior beliefs, but the opposite is true for the effects on policy index.

only on the average tax rate are denoted in red bins.[45] In the control group, a negligible share of respondents guessed that their billionaire paid a tax rate close to ProPublica's estimate. In the treatment groups, a near-majority of subjects guessed that their billionaire paid a tax rate close to ProPublica's estimate. The fact that the respondents updated their beliefs so strongly suggests that they were paying attention to the information and trusted it. Furthermore, the provision of information about the tax rate shifted beliefs to the left: relative to the control group, receiving ProPublica's estimate lowered the perceived tax rate by about 5.5 pp (p-value<0.001). In turn, receiving additional information about tax loopholes did not have any *additional* effect on the perceived tax rate (difference p-value = 0.322), as shown in panel (b) of Figure D.12.

Evidence from the emotions survey suggests that in addition to finding the information trustworthy, most subjects had an emotional reaction to it. While there is a small share of positive reactions to the tax rate message (11.8%, from Figure 4), there is a substantial share of negative emotions (83.8%). The analysis of the open-ended responses suggests the most common reactions were a general sense of unfairness and injustice (41.2% of subjects, from Table D.5) and anger and frustration due to the wealthy not paying their fair share in taxes (33.3%).

The last two rows of Table 1 report the estimates for the average treatment effects in the tax rate treatment arm. There are two coefficients. The Tax Rate coefficient corresponds to the effects of the tax rate information. The Tax Loophole coefficient corresponds to the effect of showing additional information on tax loopholes. The tax rate treatment causes a stronger perception that billionaires' taxes are unfairly low: column (3) shows a positive effect of 0.327 points (p-value<0.001), or about 0.221 standard deviations. Given that the tax rate treatment persuaded subjects that billionaires pay lower tax rates and that they pay unfairly low taxes, one may expect this treatment to raise the demand for taxation. In contrast, we find a robust negative effect. Column (1) shows that the tax rate information has a negative effect on the desired income tax rate of about 3.9 pp (p-value<0.001); column (2) shows a negative effect on the corporate tax rate of 3.9 pp (p-value<0.001); and column (4) shows a negative effect on policy support by 0.112 standard deviations (p-value=0.001). All of these negative effects are strongly significant even after accounting for multiple hypothesis testing (q-values of <0.001, <0.001 and 0.006, respectively).

Our preferred interpretation is that these negative effects reflect a reference-point or status quo bias that we term *tax-hike aversion*. We define tax-hike aversion as a reference-dependent reluctance to choose a tax rate substantially above the rate an individual believes is currently

[45] Additional comparisons for the treatment group who received additional information on tax loopholes are shown in panel (a) of Figure D.12.

in place. Thus, when information lowers the perceived status quo tax rate, respondents may prefer a more modest increase even as they come to view the current rate as more unfair. Appendix A provides a simple model that formalizes this argument. Moreover, the model yields two predictions that we can test with our experimental data: (i) the negative effect of the tax rate treatment on desired tax rates should be concentrated among individuals who initially overestimated billionaires' tax rates (see Proposition A.1); and (ii) similarly, the positive effect on perceived unfairness should be concentrated among those who overestimated billionaires' tax rates (see Proposition A.2).

We can illustrate the intuition behind these two predictions with a simple example. Consider an individual who believes that ideally (i.e., in the absence of a status-quo bias) billionaires should pay a 40% tax rate and also believes that they currently pay 40%. If assigned to the control group, this individual would choose a billionaire tax rate of 40%. If assigned to the treatment group, however, the individual would learn that billionaires actually pay 16%. In this case, raising the tax rate from 16% all the way to 40% would feel excessive due to tax-hike aversion. Rather than fully adjusting to the ideal rate, the individual would "settle" on a rate somewhere between 16% and 40%. In other words, the treatment would reduce the chosen tax rate relative to the control condition. Regarding perceived unfairness, if assigned to the control group, the individual would believe that billionaires should pay 40% and that they in fact pay 40%, and would therefore perceive the tax as fair. If assigned to the treatment group, by contrast, the individual would believe that while billionaires should ideally pay a 40% tax, they pay only 16%, leading taxes to be perceived as unfairly low.

To test these two predictions, we turn to the heterogeneity analysis by prior beliefs. The results are presented in Appendix D.9. Consistent with the tax-hike aversion channel, the negative effects of the tax rate treatment on desired tax rates and the positive effects on perceived unfairness are concentrated among respondents who initially overestimated billionaires' tax rates. An exception is policy support, for which the negative effect is larger among respondents who initially underestimated billionaires' tax rates. One possible interpretation is that respondents who learn that billionaires pay higher rates than they previously believed perceive less need for additional policies aimed at increasing their taxation. Because this pattern does not correspond to the clean prediction of the tax-hike-aversion model, it also reinforces the need to interpret tax-hike aversion as one possible mechanism rather than as the sole explanation. Moreover, our interpretation of tax-hike aversion is consistent with evidence from Charite et al. (2022). In a laboratory experiment in which individuals can redistribute resources between third parties, the authors document reference-point effects around the status quo allocation. More generally, Rees-Jones (2017) provides evidence of reference points in the context of taxation, showing that loss aversion affects taxpayers as

they file their annual tax returns. Together, these patterns provide evidence consistent with tax-hike aversion and motivate it as our preferred interpretation.

Although tax-hike aversion is our preferred interpretation, it is not the only possible explanation for these findings. We next discuss two alternatives. First, after learning that billionaires pay a surprisingly low effective tax rate, respondents might infer that billionaires are more capable of avoiding or evading taxes than they previously believed. More broadly, respondents might apply a form of Laffer-curve logic, reasoning that higher statutory rates would induce greater avoidance or reduce taxable activity. They may also expect geographically mobile billionaires to relocate across states or countries in response to higher rates, thereby limiting the additional revenue raised. Indeed, experimental evidence suggests that individuals take the mobility of high-income earners into account when choosing tax rates (Engelmann et al., 2023). However, some of our evidence weighs against avoidance and enforcement concerns being the sole explanation. The tax rate information does not increase respondents' belief that billionaires abuse tax loopholes. Moreover, as discussed below, explicitly providing information about tax loopholes does not reduce desired tax rates. The tax rate treatment also reduces support for additional IRS funding, whereas a desire to address avoidance through stronger enforcement might have produced the opposite effect. These findings do not rule out behavioral responses or mobility concerns, but they suggest that the avoidance-and-enforcement channel cannot fully account for the results.

A second possibility is a backfiring channel. Some respondents may distrust the tax rate estimate or perceive its source as politically motivated advocacy for raising taxes on billionaires. If this perception generates reactance or counterarguing, exposure to the message could reduce support for the position it is perceived to promote, thereby lowering desired tax rates. Simple distrust in the information would not by itself generate this pattern: it would instead weaken belief updating and attenuate the treatment effect toward zero. A negative effect therefore requires distrust to trigger a backfiring response. The available evidence provides little support for this mechanism as the primary explanation. On average, subjects update their beliefs substantially toward the information provided, suggesting that they treated it as credible. Moreover, the treatment effects do not differ meaningfully by political affiliation, as might be expected if partisan reactions to the source were central. Although these patterns do not rule out a backfiring effect completely, they suggest that it is unlikely to account for most of the negative treatment effect.

Next, we discuss the effects of the additional information about the tax loopholes. Although most people already believed that billionaires abuse the tax system, this information reinforces that belief even further. More precisely, panel (f) of Figure 3 shows the distribution of posterior beliefs that billionaires abuse the tax code to avoid taxes across the two treat-

ment groups. Figure D.12 provides additional comparisons with the control group and shows that beliefs are strongly right-skewed even in the control group. Even among individuals who did not receive any information on the tax rate or tax loophole, a majority (57.6%) already strongly agreed (i.e., chose a 6 on a 0–6 scale) with the statement that billionaires abuse the tax code. Information about tax rates on its own does not have a significant impact on this belief (p-value=0.471), as shown in panel (c) of Figure D.12. However, the additional information about tax loopholes increases this belief by 0.20 points (p-value=0.011). The modest size of this effect may be due to some inattention: compared to the others, this message is wordier and more technical, so some individuals may not have paid close enough attention or fully understood it. However, the modest size of the effect is probably due to the very high baseline level of belief: i.e., a majority of individuals already chose the highest score in their prior beliefs, so they maxed out on the scale, leaving no room to further increase their beliefs.[46]

The data from the emotions survey indicate that reactions to the loopholes message were as negative, if not more so, than reactions to the tax rate information alone (for details, see Figure 4). This is also reflected in the open-ended data, according to which the most common reactions were anger and frustration (38.5% of subjects, as shown in Table D.5) and criticism of the unfair tax system (36.5%).

The last row of Table 1 shows the average treatment effects of the tax loophole information on the different outcomes. On the one hand, the tax loophole treatment does not appear to significantly increase the preferred tax rates: column (1) shows an effect of 0.456 (p-value=0.563) on the preferred income tax rate; and column (2) shows an effect of 0.340 (p-value=0.662) on the preferred corporate tax rate. Column (3) shows no significant effect on the belief that billionaires pay unfairly low taxes either (coefficient of 0.038, p-value=0.482). On the other hand, we observe a significant effect on policy support: column (4) shows an effect of 0.079 standard deviations (p-value=0.022, q-value=0.068).[47] We speculate that, upon finding out about tax loopholes, individuals may perceive little value in raising tax rates, since billionaires can simply avoid them. Instead, they may be more inclined to address the root of the problem through tax reform. Indeed, Table D.7 shows that the positive effects on the policy index are driven primarily by support for two specific policies aimed at mitigating the impact of existing loopholes: the minimum income tax and a new international tax for large corporations.

[46] Finding out that billionaires abuse the tax code may cause negative attitudes towards them. According to the results from column (5) of Table 1, the tax loophole information had a negative but statistically insignificant effect on the non-tax attitudes index. A closer inspection reveals a stronger picture: Table D.7 shows that the treatment had strong negative effects on perceptions of trustworthiness and the deservingness of their wealth.

[47] As reported in Table D.9, we do not find any significant heterogeneity by prior beliefs.

3.6 Additional Results and Robustness Checks

In Appendix D.10, we show that the results are consistent across alternative specifications. More precisely, we show that the estimates are similar when we use a more basic set of controls or when we do not include any control variables at all. Additionally, we show that the results are not affected by dropping individuals who are outliers in terms of their prior misperceptions, nor by excluding the minority of subjects who found the survey difficult. In the baseline specification, as specified in the pre-registration, we pool the data across the five billionaires. In Appendix D.11, we provide evidence that the results do not seem to be driven by any single billionaire, insofar the estimates are consistent when we exclude one billionaire at a time.

Appendix D.13 examines whether treatment effects vary by political affiliation, income, and education. Because the experiment was conducted with an online sample that is somewhat more left-leaning, lower-income, and more educated than the U.S. population, these heterogeneity results are important for interpreting the aggregate estimates. There are also plausible reasons to expect heterogeneity along these dimensions. For example, some studies find that the effects of information on preferences for redistribution differ by political affiliation (e.g., Karadja et al., 2017; Alesina et al., 2018; Fehr et al., 2022). In our setting, there are also reasons why the effects might vary by income or education. Higher-income individuals may interpret the tax rate information as relevant to their own future tax exposure, while less-educated individuals may be less informed initially and therefore update their beliefs more strongly in response to the information. Overall, we find few statistically significant differences by political affiliation, income, or education. One notable exception concerns policy support: the negative effect of the tax rate treatment on the policy index is concentrated among higher-income respondents. More generally, these analyses have limited statistical power. Thus, the absence of statistically significant differences in most comparisons should not be interpreted as evidence of complete treatment-effect homogeneity.

As with any survey experiment, a potential concern is experimenter demand effects; that is, subjects may alter their survey responses to please the experimenter, even if their underlying views remain unchanged. We selected treatments that could plausibly increase the demand for taxation. In fact, experts expected that all the treatments would have positive effects on the preferred top income tax rate. Thus, if experimenter demand was strong, we would expect all treatments to increase support for redistribution. Contrary to this expectation, the data show varied effects. For example, while the luxury treatment increased demand for taxation, the non-merit treatment had no effects, and the tax rate treatment had strong negative effects. To attribute our findings solely to experimenter-demand effects, one would have to assume that some treatments (e.g., luxury) induced experimenter demand but others

did not, which seems unlikely.

One common way to address concerns about experimenter demand is by looking at behavior instead of survey responses. With this in mind, we included two behavioral measures of support for taxation, which are part of the policy support index: (i) the decision to split a real donation budget between World Relief and the Americans for Tax Fairness; (ii) the decision to sign the Oxfam petition. In fact, as shown in Table D.7, we find some significant effects on behavioral outcomes. More precisely, the luxury treatment increases donations to the Americans for Tax Fairness by \$11 (p-value=0.010), and the tax rate treatment reduces the share of respondents signing the petition by 4.9 pp (p-value=0.021).

Another common way to assess whether the effects are due to experimenter demand is to look at the persistence of the effects (Cavallo et al., 2017; Haaland et al., 2023). One caveat though is that as time passes, one naturally expects the effects to dissipate, as individuals forget about the information they received in the past and may gather new data. For example, Cavallo et al. (2017) conducted an experiment that provided information on inflation to households. When they re-interviewed these households four months later, they found that 45.6% of the belief updates during the baseline survey persisted in the follow-up period. In the case of inflation, it may be expected that individuals retain valuable information more readily because it is useful on a daily basis. However, information about billionaires is not particularly actionable outside of specific contexts, such as voting in a presidential election, and for that reason, individuals may forget it more easily.

Responses to the follow-up survey show that the effects on beliefs persisted, at least partially, one month later. Figure D.13 mirrors the structure of Figure 3 and presents the distributions of posterior beliefs at follow-up.[48] Respondents were unlikely to recall the precise numerical information; however, the treatments had a persistent shifting effect on the distribution of posterior beliefs, although quantitatively weaker than at baseline.[49] For example, the luxury treatment increased the median perceived value of the billionaire's home by \$9.00 million at baseline ($p < 0.001$) and by \$8.00 million a month later in the follow-up survey ($p = 0.008$).

Table 2 shows the ATEs on the different outcomes, with a side-by-side comparison between the outcomes measured in the baseline versus the follow-up survey. Due to the attrition rate, the sample sizes are 18% smaller for follow-up outcomes than for the baseline outcomes, and thus the effects are less precisely estimated. Moreover, as documented in previous studies (e.g. Cavallo et al., 2017), and consistent with the decay of effects in posterior beliefs, we expect

[48] Additional comparisons between treatment and control groups are reported in Figure D.14.

[49] The persistence of the tax rate treatment effects also weighs against anchoring as the primary mechanism. Anchoring effects tend to be short-lived (Cavallo et al., 2017), whereas the effects on preferred income and corporate tax rates persist almost entirely one month later.

the effects on the follow-up outcomes to be a fraction of the corresponding effects on baseline outcomes. The combination of lower sample sizes and smaller effect sizes makes it more difficult to detect effects on specific follow-up outcomes. In terms of point estimates, Table 2 shows that, due to a lack of power, we often fail to reject the null hypothesis of zero effects. However, in most cases, we cannot reject the hypothesis that the point estimates are the same in the baseline and follow-up surveys. The pairwise comparisons from Table 2 suggest that the effects on follow-up outcomes are typically half as large as the corresponding effects on the baseline outcomes. The main exception is the earnings treatment, where we observe that the effects on the desired income and corporate tax rates seem to have completely disappeared after a month. As described in Section 3.4 above, the effects of the earnings treatment were noted to be weaker, as they did not affect perceived tax fairness or policy support. The lack of persistence in these effects reinforces this conclusion.

In order to maximize statistical power, rather than focus on a specific coefficient, Figure 5 provides a more systematic comparison, pooling all the different coefficients. In this scatterplot, each observation corresponds to a pair of coefficients from Table 2: the x-axis shows the average effect of a given treatment on the baseline outcome (e.g., the effect of the luxury treatment on the preferred income tax rate), while the y-axis shows the corresponding effect on the follow-up outcome.[50] To be able to compare the results between different outcomes, we standardize all coefficients using the standard deviation of the relevant outcome variable in the control group. Figure 5 shows a strong linear relationship between the effects on the baseline outcomes and the corresponding effects on the follow-up outcomes. The slope of this linear relationship suggests that, on average, the effects on follow-up outcomes were 51.3% as strong as the effects on baseline outcomes.

Given the differential follow-up participation within the tax arm, we conduct two additional checks for selective attrition (see Appendix D.12 for more details). First, following Duflo et al. (2015), Appendix Table D.13 compares treatment effects on baseline outcomes estimated in the full baseline sample with the corresponding effects estimated after restricting the baseline sample to respondents who participated in the follow-up survey. The estimates are similar across samples, and none of the differences is statistically significant at the 5% level. Second, Appendix Table D.14 reports Lee bounds for the treatment effects on follow-up outcomes. The bounds further indicate that the persistent negative effects of the tax rate treatment on preferred income and corporate tax rates are unlikely to be driven by differential attrition.

Taken together, the results reveal distinct patterns across treatments. All five information

[50] Figure D.23 shows that the results are robust when the baseline estimates are restricted to respondents who participated in both surveys and when treatment effects are compared using individual outcomes that were elicited identically at baseline and follow-up.

treatments shifted their targeted beliefs and elicited emotional reactions, but these responses did not translate mechanically into tax preferences. The luxury treatment increased preferred tax rates, perceived tax unfairness, and policy support without changing non-tax attitudes toward the billionaire. Its effects on perceived unfairness and policy support persisted at follow-up, while its effect on preferred tax rates weakened. The non-merit treatment worsened non-tax attitudes, and this effect persisted, but it did not affect preferred tax rates or policy support in either survey. The earnings treatment had little effect on perceived unfairness, policy support, or non-tax attitudes. Although it initially increased preferred tax rates, this effect disappeared at follow-up. The tax rate treatment increased perceived unfairness while reducing preferred tax rates and policy support. The negative effect on preferred tax rates remained substantial at follow-up, whereas the effects on perceived unfairness and policy support weakened. Finally, the tax loophole treatment increased beliefs about loophole abuse and initially raised policy support, but it did not increase preferred tax rates, and its effect on policy support did not persist.

Last but not least, a noteworthy caveat with the research design is that three of the treatments—non-merit, luxury, and earnings—convey information tailored to a specific billionaire, while other outcomes—such as preferred tax rates on billionaires—apply to billionaires more generally. If respondents do not extrapolate from the specific billionaire described in the treatment to the broader population of billionaires, estimated treatment effects on these broader outcomes will be attenuated. For example, it is possible that the luxury treatment would have had a stronger effect if it had included information about the lifestyles of multiple billionaires rather than focusing on a single individual. Likewise, because the non-merit treatment included information about a specific billionaire, one interpretation of the null effect is that the message was not sufficiently general to shift views about billionaires as a group.

Although this caveat is important, we believe the evidence mitigates the concern. For example, in the luxury treatment—which provides information about the home of a specific billionaire—we observe significant effects not only on billionaire-specific outcomes—such as perceptions that the taxes paid by the specific billionaire are unfair—but also on outcomes pertaining to billionaires in general—such as income and corporate tax rates and policy support. Conversely, in the tax rate treatment—which provides information about the average tax rates paid by billionaires in general—we find significant effects on both outcomes pertaining to billionaires in general but also outcomes about the specific billionaire.[51] Taken together, these patterns suggest that respondents do not learn from the information treatments in a compartmentalized manner. Information about a particular billionaire appears

[51] For a more detailed discussion, see Appendix D.8.

to shape views about billionaires more broadly, and general information about billionaires appears to influence evaluations of specific individuals.

4 Conclusions

This paper examined public demand for taxing billionaires using a pre-registered survey experiment. We provided respondents with four different information treatments about the lives and careers of well-known billionaires. We find that the public holds substantial misperceptions about key facts—for example, how much billionaires earn or pay in taxes. All four treatments significantly affected individuals' beliefs; for instance, the treatment about how billionaires accumulated their wealth led respondents to view honest and hard work as a less important factor. Across all four treatments, respondents also exhibited strong emotional reactions. Contrary to expert predictions that all four treatments would increase support for taxing billionaires, we find that only one—the treatment highlighting lavish consumption—had a robust positive effect. Moreover, one treatment—revealing that billionaires pay low effective tax rates—had a significant negative effect.

Our findings raise some questions about conventional wisdom in the academic literature. First, they suggest that classic fairness distinctions—such as the difference between outcomes due to luck versus merit—that feature prominently in laboratory settings do not directly translate into support for taxing billionaires, perhaps because people view real-world economic success as stemming from a complex mix of factors. Second, our results highlight the limits of factual or statistical information: quantitative information appeared to have less impact than qualitative narratives and imagery. This comparison should be interpreted cautiously, however, because the luxury treatment changes the content and format of the information simultaneously and therefore does not provide a clean test of qualitative versus quantitative communication. Third, we provide evidence consistent with a status quo or reference-point effect: when individuals learn that billionaires currently pay very low tax rates, they become less inclined to support large increases.

A key consideration in any empirical study is the external validity of its findings. Ultimately, this is an empirical question, and future research will help determine the extent to which our results generalize to other contexts. In the terminology of List (2020), our findings represent a wave-1 insight: they provide initial causal evidence and a first test of the theory. Nevertheless, we offer some guidance on the contextual features that should be considered before extrapolating these results. Because our analysis focuses on the United States, an important question is how these results might generalize to other countries. Certain features of the U.S. context could make the effects weaker or stronger elsewhere. For example, the

United States exhibits comparatively high levels of inequality acceptance relative to other countries (Almås et al., 2025, 2020). In countries where demand for redistribution is already high, there may be less scope for information about billionaires to further increase support for taxing them. Conversely, billionaires in other countries may receive less media attention than in the United States, either because public interest is lower or because they maintain a more discreet public profile. In such settings, providing information about their wealth and lifestyles could have stronger effects, as this information may be more novel to the public.

Another consideration is that, in most of our treatments, we provided information about specific billionaires who are well known and highly visible in the media. This was an intentional design choice, as it allowed us to leverage their public profiles to construct treatments that were concrete, factual, and relatable. Although we show that our results are not driven by any single prominent billionaire, it remains an open question how the effects might differ if the information concerned less prominent figures. The focus on highly visible billionaires does not yield a clear upper- or lower-bound interpretation of our estimated effects. Respondents are likely to know less about less prominent billionaires, leaving more room for new information to change their beliefs. At the same time, information about less prominent individuals may be less salient or less relevant to broader tax preferences, weakening the translation of belief updating into policy demand. Because these forces operate in opposite directions, we cannot determine whether the corresponding effects for less prominent billionaires would be larger or smaller.

Finally, we acknowledge that our treatments do not capture all dimensions of billionaires' public image that may shape attitudes toward taxation. Public attitudes may depend, for instance, on whether billionaires are associated with philanthropy or with significant product innovations. For example, figures such as Bill Gates are often viewed through the lens of their philanthropic work, whereas others, such as Steve Jobs, are primarily associated with influential products like those developed at Apple. All of these additional considerations—how the effects may vary across countries, among less prominent billionaires, and across other dimensions of billionaires' public image—represent avenues for future research.

Our findings also have implications for the design of communication strategies around tax policy. Our results highlight the importance of carefully considering how such information is communicated to the public. For example, recent efforts have focused on estimating and publicizing the low effective tax rates that billionaires face (e.g., Leiserson and Yagan, 2021; EU Tax Observatory, 2024; Balkir et al., 2025). Evidence from the tax rate treatment shows that this information can have unexpected negative effects, potentially due to tax-hike aversion. Our findings may also inform politicians' communication strategies. For example, some politicians have explicitly referred to the immense wealth of billionaires or to their

alleged immoral behavior to build support for taxing the ultra-rich (e.g., The Guardian, 2026). However, our results from the earnings and non-merit treatments indicate that such messages have, at most, weak effects on policy preferences. By contrast, the findings from the luxury treatment suggest that communication strategies highlighting the conspicuous or lavish lifestyles of billionaires (e.g., The Guardian, 2025) may be more effective in increasing support for higher taxes. Whether these communication effects are specific to billionaires or extend to wealthy individuals more broadly remains an open question.

References

Alesina, Alberto, Armando Miano, and Stefanie Stantcheva (2023) "Immigration and redistribution," *The Review of Economic Studies*, 90 (1), 1–39.

Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso (2018) "Intergenerational mobility and preferences for redistribution," *American Economic Review*, 108 (2), 521–54.

Almås, Ingvild, Alexander W Cappelen, Erik Ø Sørensen, and Bertil Tungodden (2025) "Fairness across the world," *Working Paper*.

Almås, Ingvild, Alexander W Cappelen, and Bertil Tungodden (2020) "Cutthroat capitalism versus cuddly socialism: Are Americans more meritocratic and efficiency-seeking than Scandinavians?" *Journal of Political Economy*, 128 (5), 1753–1788.

Anderson, Michael L (2008) "Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects," *Journal of the American Statistical Association*, 103 (484), 1481–1495.

Andre, Peter (2024) "Shallow Meritocracy," *The Review of Economic Studies, forthcoming*.

Balkir, Akcan S, Emmanuel Saez, Danny Yagan, and Gabriel Zucman (2025) "How Much Tax Do US Billionaires Pay? Evidence from Administrative Data," *NBER Working Paper*.

Benjamini, Yoav, Abba M Krieger, and Daniel Yekutieli (2006) "Adaptive linear step-up procedures that control the false discovery rate," *Biometrika*, 93 (3), 491–507.

Bortolotti, Stefania, Ivan Soraperra, Matthias Sutter, and Claudia Zoller (2023) "Too Lucky To Be True: Fairness Views under the Shadow of Cheating," *The Review of Economics and Statistics, forthcoming*.

de Bresser, Jochem and Marike Knoef (2022) "Eliciting preferences for income redistribution: A new survey item," *Journal of Public Economics*, 214, 104724.

Cappelen, Alexander W, Astri Drange Hole, Erik Ø Sørensen, and Bertil Tungodden (2007) "The pluralism of fairness ideals: An experimental approach," *American Economic Review*, 97 (3), 818–827.

Cappelen, Alexander W., Morten N. Støstad, and Bertil Tungodden (2025) "Majority in 40 Countries Support Unprecedented Coordinated Billionaire Tax Discussed by G20," *Working Paper*.

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

CBS (2019) "Interview with Alexandria Ocasio-Cortez," *Interview with 60 Minutes, aired on January 6, 2019*.

Charite, Jimmy, Raymond Fisman, Ilyana Kuziemko, and Kewei Zhang (2022) "Reference points and redistributive preferences: Experimental evidence," *Journal of Public Economics*, 216, 104761.

Cohn, Alain, Lasse J Jessen, Marko Klašnja, and Paul Smeets (2023) "Wealthy Americans and redistribution: The role of fairness preferences," *Journal of Public Economics*, 225, 104977.

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

Del Valle, Gaby (2018) "The appealing myth of the frugal billionaire," *Vox*, https://www.vox.com/the-goods/2018/9/27/17907312/warren-buffett-mark-zuckerberg-mark-cuban-frugal-billionaires, Accessed: 2025-06-16.

DellaVigna, Stefano, Nicholas Otis, and Eva Vivalt (2020) "Forecasting the Results of Experiments: Piloting an Elicitation Strategy," *AEA Papers and Proceedings*, 110, 75–79.

Di Tella, Rafael, Juan Dubra, and Alejandro Lagomarsino (2021) "Meet the Oligarchs: Business Legitimacy and Taxation at the Top," *The Journal of Law and Economics*, 64 (4), 651–674.

Di Tella, Rafael, Ricardo Perez-Truglia, Andres Babino, and Mariano Sigman (2015) "Conveniently Upset: Avoiding Altruism by Distorting Beliefs about Others' Altruism," *American Economic Review*, 105 (11), 3416–42.

Dohmen, Thomas and Tomáš Jagelka (2024) "Accounting for individual-specific reliability of self-assessed measures of economic preferences and personality traits," *Journal of Political Economy Microeconomics*, 2 (3), 399–462.

Dreber, Anna, Thomas Pfeiffer, Johan Almenberg, Siri Isaksson, Brad Wilson, Yiling Chen, Brian A Nosek, and Magnus Johannesson (2015) "Using prediction markets to estimate the reproducibility of scientific research," *Proceedings of the National Academy of Sciences*, 112 (50), 15343–15347.

Duflo, Esther, Pascaline Dupas, and Michael Kremer (2015) "Education, HIV, and early fertility: Experimental evidence from Kenya," *American Economic Review*, 105 (9), 2757–2797.

Durante, Ruben, Louis Putterman, and Joël Van der Weele (2014) "Preferences for redistribution and perception of fairness: An experimental study," *Journal of the European Economic Association*, 12 (4), 1059–1086.

Elías, Julio J, Nicola Lacetera, and Mario Macis (2019) "Paying for kidneys? A randomized survey and choice experiment," *American Economic Review*, 109 (8), 2855–2888.

Engelmann, Dirk, Eckhard Janeba, Lydia Mechtenberg, and Nils Wehrhöfer (2023) "Preferences over taxation of high-income individuals: Evidence from a survey experiment," *European Economic Review*, 157, 104505.

EU Tax Observatory (2024) "Global Tax Evasion Report 2024," https://www.taxobservatory.eu//www-site/uploads/2023/10/global_tax_evasion_report_24.pdf.

Fehr, Dietmar, Johanna Mollerstrom, and Ricardo Perez-Truglia (2022) "Your Place in the World: Relative Income and Global Inequality," *American Economic Journal: Economic Policy*, 14 (4), 232–68.

Fehr, Ernst, Thomas Epper, and Julien Senn (2026) "Social preferences and redistributive politics," *Review of Economics and Statistics*, 1–16.

Fisman, Raymond, Keith Gladstone, Ilyana Kuziemko, and Suresh Naidu (2020) "Do Americans want to tax wealth? Evidence from online surveys," *Journal of Public Economics*, 188, 104207.

Haaland, Ingar, Christopher Roth, Stefanie Stantcheva, and Johannes Wohlfart (2024) "Measuring what is top of mind," *Harvard University Working Paper*.

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart (2023) "Designing information provision experiments," *Journal of Economic Literature*, 61 (1), 3–40.

Hansen, Kristina (2023) "Greed, Envy, and Admiration: The Distinct Nature of Public Opinion about Redistribution from the Rich," *American Political Science Review*, 117 (1), 217â234.

Hauser, Oliver P and Michael I Norton (2017) "(Mis)perceptions of Inequality," *Current Opinion in Psychology*, 18, 21–25, https://doi.org/10.1016/j.copsyc.2017.07.024.

Hope, David, Julian Limberg, and Nina Weber (2023) "Why do (some) ordinary Americans support tax cuts for the rich? Evidence from a randomised survey experiment," *European Journal of Political Economy*, 78, 102349.

Kane, John V and Benjamin J Newman (2023) "What they have but also who they are: Avarice, elitism, and public support for taxing the rich," *Public Opinion Quarterly*, 87 (2), 402–414.

Karadja, Mounir, Johanna Mollerstrom, and David Seim (2017) "Richer (and holier) than thou? The effect of relative income improvements on demand for redistribution," *Review of Economics and Statistics*, 99 (2), 201–212.

Kiatpongsan, Sorapop and Michael I. Norton (2014) "How Much (More) Should CEOs Make? A

Universal Desire for More Equal Pay," *Perspectives on Psychological Science*, 9 (6), 587–593.

Krueger, Alan B. and David A. Schkade (2008) "The reliability of subjective well-being measures," *Journal of Public Economics*, 92 (8), 1833–1845.

Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva (2015) "How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments," *American Economic Review*, 105 (4), 1478–1508.

Leiserson, Greg and Danny 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.*

List, John A (2020) "Non est disputandum de generalizability? A glimpse into the external validity trial," *NBER Working Paper*.

Manski, Charles F and Francesca Molinari (2010) "Rounding probabilistic expectations in surveys," *Journal of Business & Economic Statistics*, 28 (2), 219–231.

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

Nathan, Brad, Ricardo Perez-Truglia, and Alejandro Zentner (2023) "Paying Your Fair Share: Perceived Fairness and Tax Compliance.," *NBER Working Paper No. 32588*.

New York Times (2019) "What Could Raising Taxes on the 1% Do?" *October 6, 2019*.

ProPublica (2023) "America's Highest Earners And Their Taxes Revealed," https://projects.propublica.org/americas-highest-incomes-and-taxes-revealed/.

Ragusa, Jordan Michael (2015) "Socioeconomic stereotypes: Explaining variation in preferences for taxing the rich," *American Politics Research*, 43 (2), 327–359.

Rasooly, Itzhak (2024) "Do the Numbers Matter? An Experiment on Policy Preferences," *Working Paper*.

Rees-Jones, Alex (2017) "Quantifying Loss-Averse Tax Manipulation," *The Review of Economic Studies*, 85 (2), 1251–1278.

Rosen, Sherwin (1981) "The economics of superstars," *The American Economic Review*, 71 (5), 845–858.

Saez, Emmanuel and Gabriel Zucman (2020) "The rise of income and wealth inequality in America: Evidence from distributional macroeconomic accounts," *Journal of Economic Perspectives*, 34 (4), 3–26.

Stantcheva, Stefanie (2021) "Understanding tax policy: How do people reason?" *The Quarterly Journal of Economics*, 136 (4), 2309–2369.

Tax Foundation (2023) "Biden Budget 2025 Tax Proposals," https://taxfoundation.org/research/all/federal/biden-budget-2025-tax-proposals/.

The Economist (2024) "How to Tax Billionaires – and How Not To," *The Economist, June 19, 2024*.

The Guardian (2025) "The American system is badly broken," *September 24, 2025*.

——— (2026) "Bernie Sanders rails against billionaire "greed" amid California tax battle," *February 19, 2026*.

Trump, Kris-Stella (2025) "When is it Fair to Tax the Rich? The Importance of Pro-Social Behavior," *Comparative Political Studies*, 00104140241302716.

Yusof, Jeffrey ⓡ Simona Sartor (2024) "Market Luck: Skill-Biased Inequality and Redistributive Preferences," *Working paper*.

Zucman, Gabriel (2024) "Opinion: Global Billionaires Should Pay More Taxes," *The New York Times, May 3 2024*.

Figure 1: Screenshots of Treatments (Using Bill Gates as an Example)

(a) Luxury

According to some accounts, **Bill Gates lives in a $130 million mansion in Medina, Washington**. The 66,000 square-foot complex features a pool with an underwater music system, a trampoline room, a 2,500 square foot gym, a movie theater, a library, and a reception hall that can accommodate up to 200 guests.

Source: Business Insider

(b) Non-Merit

Bill Gates faced accusations of stealing the idea for Windows.

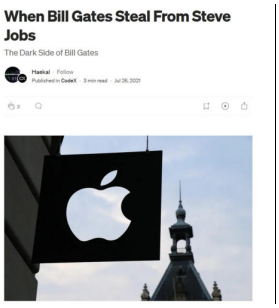

After the launch of the first version of Windows in 1985, Steve Jobs claimed that Gates had stolen the idea for the graphical user interface from Apple. "They just ripped us off completely, because Gates has no shame," Jobs once said. In response, Gates admitted that they both he and Jobs had copied the idea from the Xerox research institute: "I think it's more like we both had this rich neighbor named Xerox and I broke into his house to steal the TV set and found out that you had already stolen it."

Source: Medium, Business Insider

(c) Hourly Earnings

Bill Gates' earnings have significantly increased since he became a billionaire in 1987.

Between 1987 and 2023, his wealth has grown from around $1.25 billion to $111 billion.

This means that, on average, Bill Gates has earned about $3.05 billion per year during this period.

If we divide this yearly amount by the number of hours in a year, it means **Bill Gates has earned about $348,015 per hour** during this time.

Source: Forbes, 2023.

(d) Annual Earnings

Bill Gates' earnings have significantly increased since he became a billionaire in 1987.

Between 1987 and 2023, his wealth has grown from around $1.25 billion to $111 billion.

This means that, on average, **Bill Gates has earned about $3.05 billion per year** during this period.

Source: Forbes, 2023.

(e) Tax Rate

A single American worker who earns $45,000 per year has a total tax rate of 21% on average.

Bill Gates is one of the 25 richest Americans by net worth. Recent research has shown that **people in this group have a total tax rate of 16% on average**.

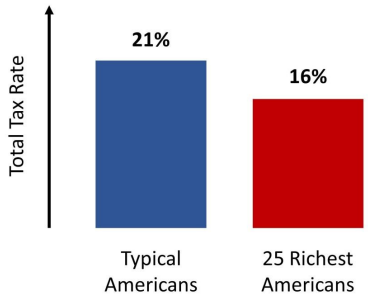

Source: ProPublica, 2022

(f) Tax Loophole

According to some accounts, billionaires manage to pay even lower tax rates through a variety of accounting strategies.

Billionaires often get paid in stocks.

Billionaires do not need to pay taxes on the stocks they hold until they sell it. By refraining from selling their stocks, billionaires can avoid generating taxable income. Instead, billionaires can access their wealth by borrowing against their stocks, which does not incur taxes.

Billionaires set up smaller companies in tax havens.

Billionaires can transfer the profits of their companies to these smaller companies, which pay very little or no taxes. A recent study revealed that U.S. multinational companies moved over half of their foreign profits to low-tax countries, leading to an estimated loss of around $50 billion in tax revenue for the US government.

Sources: ProPublica, 2021, Tørsløv, Wier, Zucman (2023)

Figure 2: A Selection of Baseline Preferences and Attitudes

(a) Preferred Top Income Tax Rate

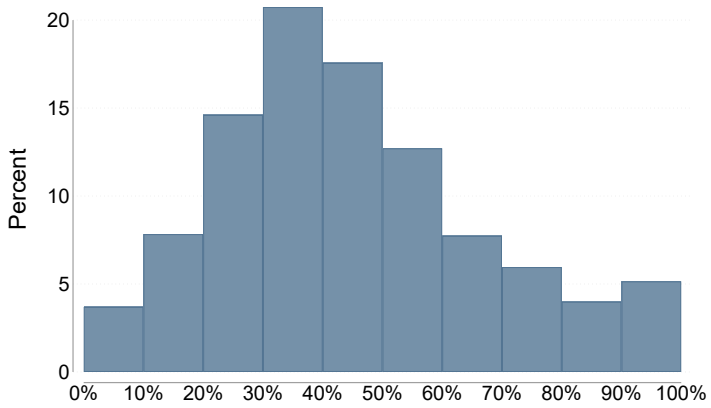

(b) Preferred Corporate Tax Rate

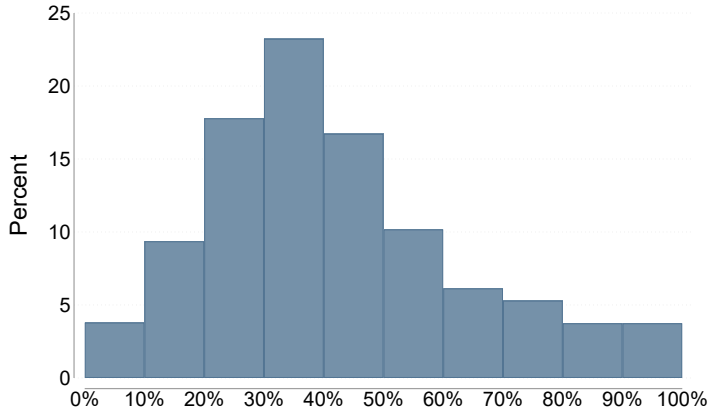

(c) Unfairness of Billionaire Taxes

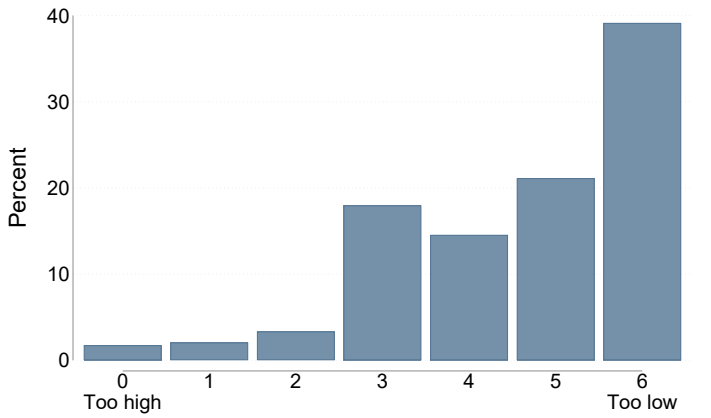

(d) Billionaire Minimum Income Tax

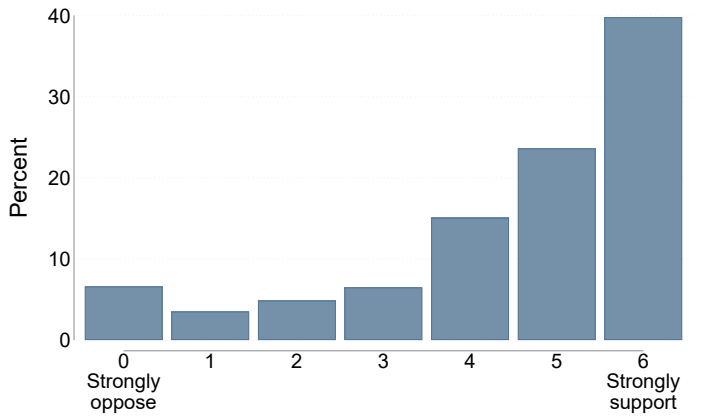

(e) Share Donated to Americans for Tax Fairness

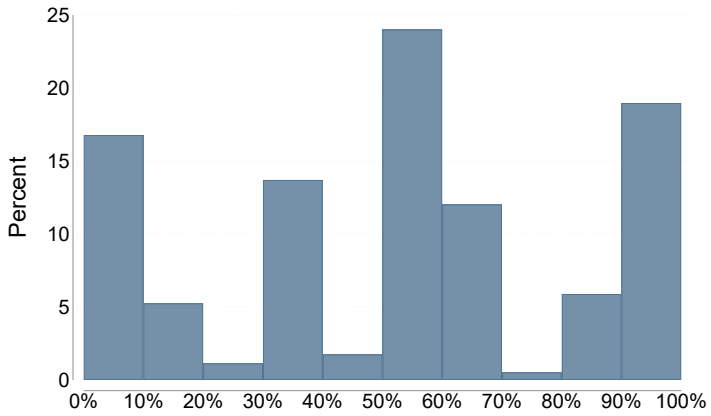

(f) Deservingness of Wealth

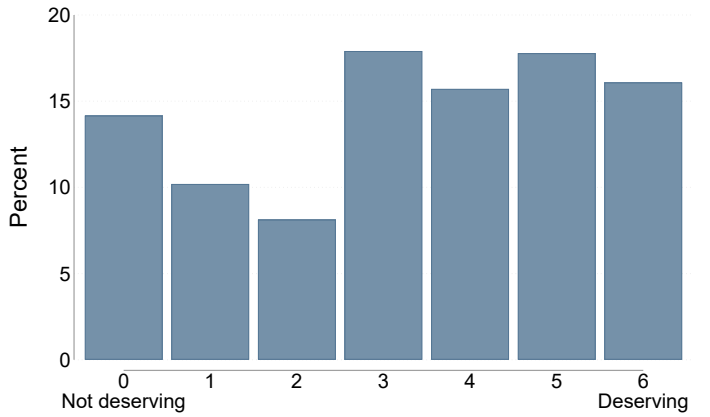

Notes: Histograms of a selection of preferences and attitudes for subjects in the control group. All survey questions used to measure these attitudes are listed in Table D.2. Panel (a) is for the preferred top income tax rate for billionaires. Panel (b) is for the preferred corporate tax rate for billionaire-owned companies. Panel (c) is for opinions on whether taxes paid by billionaires are too high or too low from a fairness perspective. Panel (d) is for the support for a policy proposal on a billionaire minimum income tax. Panel (e) is for the share of the budget donated to Americans for Tax Fairness. Panel (f) is for the opinions on whether billionaires deserve their wealth.

Figure 3: Distribution of Posterior Beliefs

(a) Perceived Value of Billionaire's Home

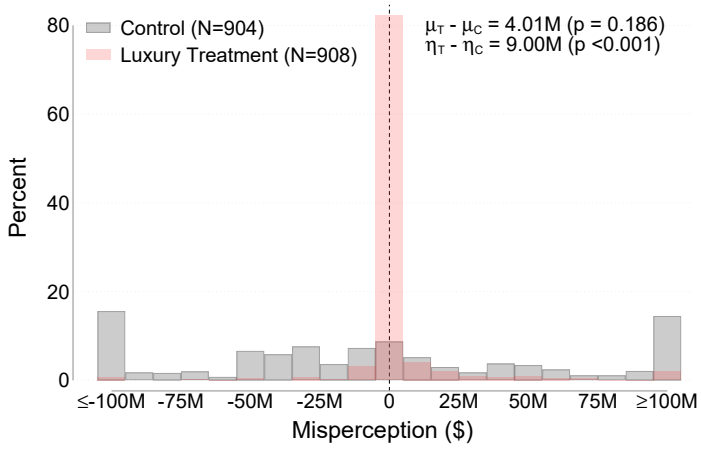

(b) Perceived Role Of Hard Work

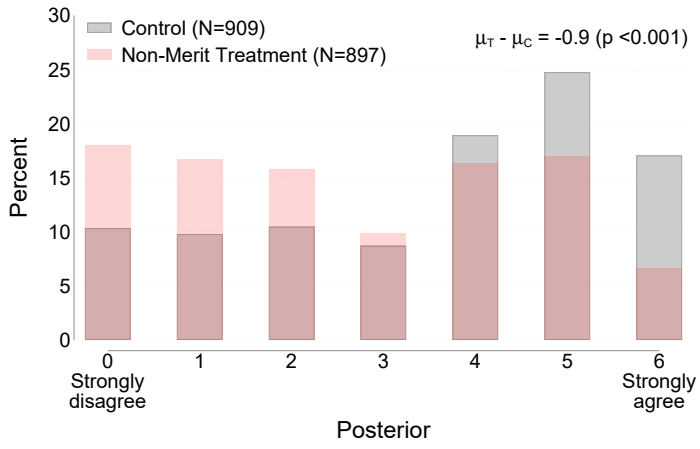

(c) Perceived Hourly Earnings

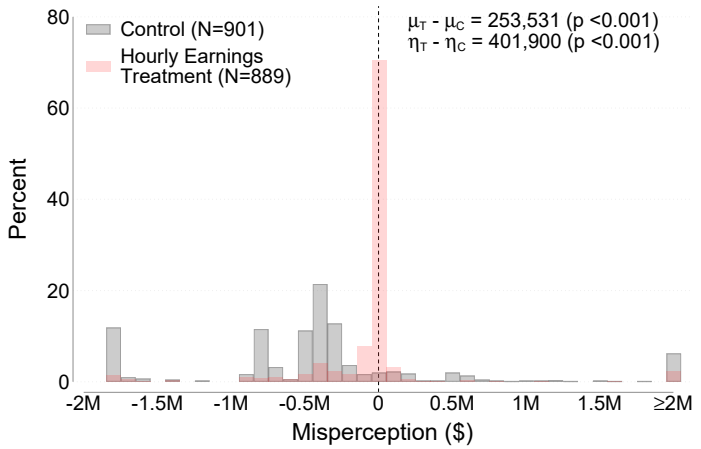

(d) Perceived Annual Earnings

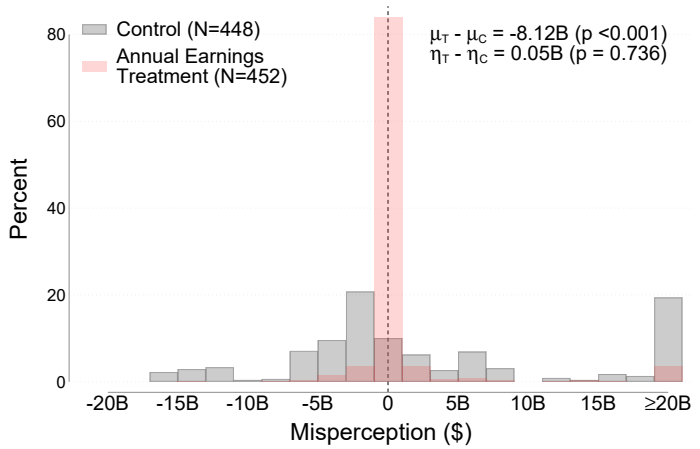

(e) Perceived Billionaire Tax Rate

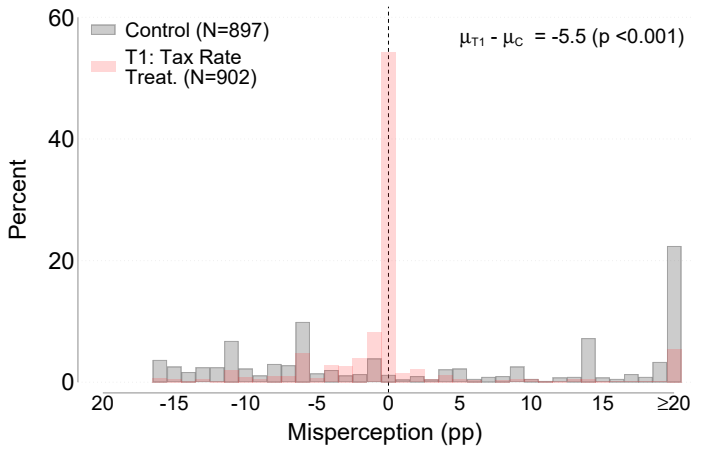

(f) Perceived Abuse of Tax Loopholes

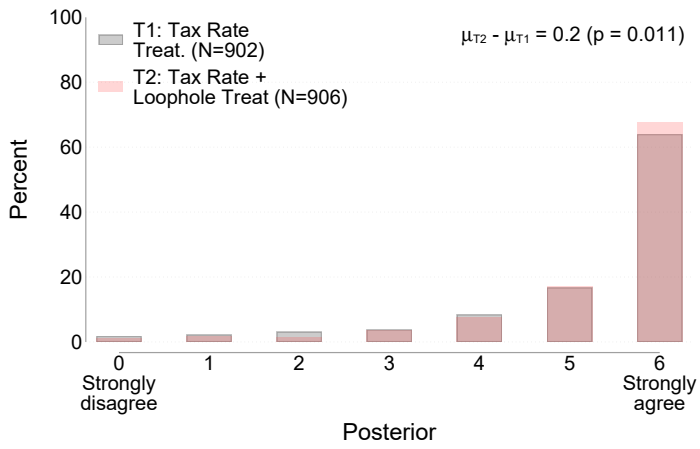

Notes: Histograms of posterior beliefs in the baseline survey. Each panel represents a different treatment arm. All survey questions used to measure these beliefs are listed in Table D.1. In panels (a), (c), (d), and (e), the x-axis corresponds to the difference between the subject's posterior belief and the information provided in the treatment. Average posteriors are denoted by μ and median posteriors by η, with the difference p-values reported in parentheses.

Figure 4: Emotions Selected in the Multiple-Choice Question from Emotions Survey

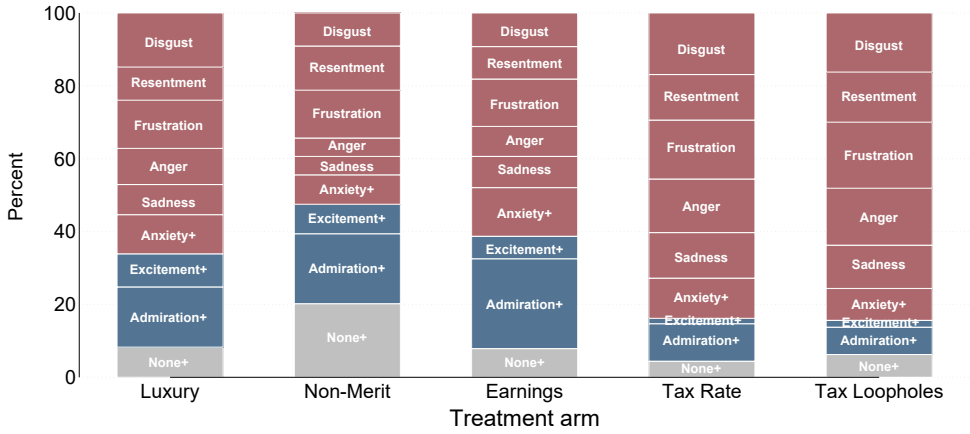

Notes: The figure shows the share of each emotion among all reported emotions within a treatment. Subjects were asked to select the emotions they experienced after reading information about the billionaire from a list of 14 emotions. Multiple emotions could be reported. Some emotions are presented as combined categories in this figure: *None+* includes the options "None" and "Indifference", *Admiration+* includes "Admiration", "Motivation", and "Inspiration", *Excitement+* includes "Excitement" and "Happiness", and *Anxiety+* includes "Anxiety" and "Envy".

Figure 5: Persistence of the Average Treatment Effects

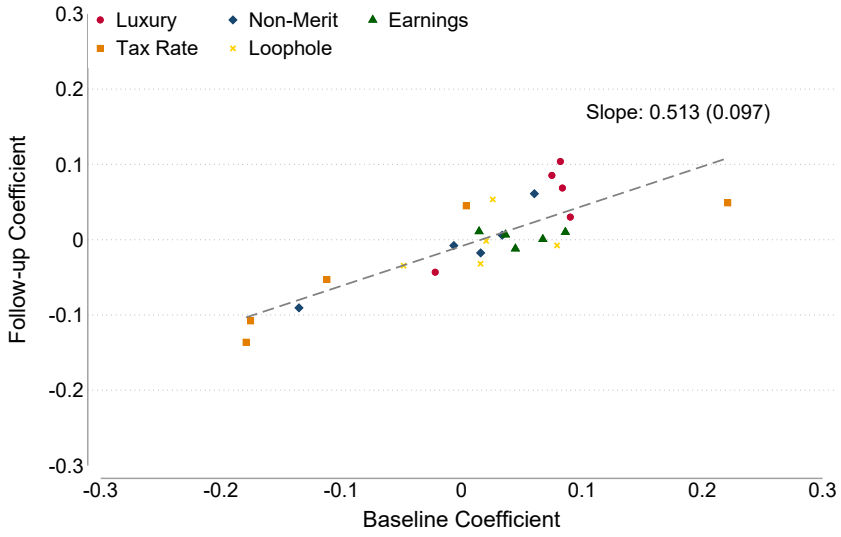

Notes: The figure compares standardized ATEs in the baseline and follow-up surveys using the estimates from Table 2. Each observation represents a treatment-outcome pair, with the baseline estimate on the x-axis and the corresponding follow-up estimate on the y-axis. Coefficients are standardized by the control-group standard deviation of the relevant outcome. Symbols and colors distinguish the different treatments.

Table 1: Average Treatment Effects

	(1) Income tax rate	(2) Corporate tax rate	(3) Taxes unfair	(4) Policy Index	(5) Non-Tax Attitudes Index	(6) N
Luxury	2.018**	1.812*	0.121**	0.075**	-0.022	1,812
	(0.961)	(0.934)	(0.061)	(0.035)	(0.040)	
	[0.083]	[0.095]	[0.094]	[0.083]	[0.476]	
Non-Merit	-0.144	0.730	0.089	0.016	-0.135***	1,806
	(1.001)	(0.967)	(0.058)	(0.035)	(0.040)	
	[0.585]	[0.467]	[0.154]	[0.476]	[0.004]	
Earnings	1.927**	1.460*	0.054	0.015	0.045	2,690
	(0.816)	(0.787)	(0.049)	(0.028)	(0.033)	
	[0.065]	[0.106]	[0.278]	[0.476]	[0.202]	
Tax rate	-3.918***	-3.875***	0.327***	-0.112***	0.004	2,705
	(0.837)	(0.833)	(0.057)	(0.035)	(0.039)	
	[0.001]	[0.001]	[0.001]	[0.006]	[0.585]	
Tax loophole	0.456	0.340	0.038	0.079**	-0.048	
	(0.788)	(0.777)	(0.054)	(0.035)	(0.039)	
	[0.476]	[0.476]	[0.474]	[0.068]	[0.235]	
Baseline Mean	42.52	39.80	4.62	0.00	0.00	
Baseline SD	22.34	21.63	1.48	1.00	1.00	

Notes: This table reports OLS estimates of equation (1) from Section 3.2 for the different treatment arms. The dependent variable in column (1) is the preferred top income tax rate for billionaires. The dependent variable in column (2) is the preferred corporate tax rate. The dependent variable in column (3) indicates whether subjects perceive billionaire taxes as too high (0) or too low (6) from a fairness perspective. The dependent variable in column (4) is an index of policy support. The dependent variable in column (5) is a non-tax attitudes index. Column (6) indicates the number of observations in each treatment arm. The baseline mean and standard deviation from the control group are presented at the bottom of the table. Significant at *10%, **5%, ***1%. Robust standard errors in parentheses. Sharpened q-values in square brackets.

Table 2: Comparison of Average Treatment Effects between Baseline and Follow-Up Surveys

	Income tax rate		Corporate tax rate		Taxes unfair		Policy Index		Non-Tax Att. Index		
	(1) Baseline	(2) Follow-up	(3) Baseline	(4) Follow-up	(5) Baseline	(6) Follow-up	(7) Baseline	(8) Follow-up	(9) Baseline	(10) Follow-up	(11) N
Luxury	2.018**	0.671	1.812*	1.485	0.121**	0.153**	0.075**	0.085**	-0.022	-0.043	1,812
	(0.961)	(1.086)	(0.934)	(1.036)	(0.061)	(0.068)	(0.035)	(0.039)	(0.040)	(0.046)	[1,482]
Diff. p-value	0.140		0.718		0.546		0.737		0.469		
Non-Merit	-0.144	-0.174	0.730	0.129	0.089	0.090	0.016	-0.018	-0.135***	-0.091**	1,806
	(1.001)	(1.128)	(0.967)	(1.081)	(0.058)	(0.064)	(0.035)	(0.038)	(0.040)	(0.046)	[1,515]
Diff. p-value	0.975		0.521		0.986		0.285		0.153		
Earnings	1.927**	0.225	1.460*	0.013	0.054	0.010	0.015	0.011	0.045	-0.012	2,690
	(0.816)	(0.901)	(0.787)	(0.870)	(0.049)	(0.052)	(0.028)	(0.032)	(0.033)	(0.037)	[2,232]
Diff. p-value	0.029		0.067		0.324		0.891		0.021		
Tax rate	-3.918***	-2.410**	-3.875***	-2.944***	0.327***	0.072	-0.112***	-0.053	0.004	0.045	2,705
	(0.837)	(0.961)	(0.833)	(0.920)	(0.057)	(0.066)	(0.035)	(0.041)	(0.039)	(0.044)	[2,201]
Diff. p-value	0.045		0.233		<0.001		0.078		0.172		
Tax loophole	0.456	-0.036	0.340	-0.691	0.038	0.079	0.079**	-0.008	-0.048	-0.034	
	(0.788)	(0.915)	(0.777)	(0.872)	(0.054)	(0.064)	(0.035)	(0.041)	(0.039)	(0.044)	
Diff. p-value	0.508		0.182		0.445		0.010		0.653		
Control Mean	42.52	42.18	39.80	39.30	4.62	4.61	0.00	0.00	0.00	0.00	
Control SD	22.34	22.78	21.63	21.89	1.48	1.48	1.00	1.00	1.00	1.00	

Notes: Results from equation (1) from Section 3.2. We estimate two regressions per treatment arm (one for the baseline survey and another for the follow-up). The dependent variables are: the preferred top income tax rate for billionaires in columns (1)–(2); the preferred corporate tax rate in columns (3)–(4); whether subjects perceive billionaire taxes as too high from a perspective of fairness in columns (5)–(6); the index of policy support in columns (7)–(8); and the non-tax attitudes index in columns (9)–(10). Column (11) indicates the number of observations in each treatment arm, with the number of observations in the follow-up study inside square brackets. We report the p-value from a test of equality between the baseline and follow-up coefficients. The baseline means and standard deviations are presented at the bottom of the table. Significant at *10%, **5%, ***1%. Robust standard errors in parentheses.