The Role of Tax Preparers in Individual Tax Optimization

Hadar Avivi[*] Katarzyna Bilicka[†] Jakob Brounstein[‡]
Felipe Lobel[§] Alexander Yuskavage[¶]

March 20, 2025

Abstract

We examine the role of paid tax preparers in tax optimization of individual taxpayers. First, using the universe of annual personal income tax returns in the US from 2011-2020, we provide novel descriptive evidence characterizing the paid tax preparers and the users of their services. Second, we develop new measures of tax optimization and quantify the effects of paid preparers on tax optimization, leveraging the rotation of clients among preparers. Our findings suggest that one standard deviation better tax preparers reduce the effective tax rates of their clients by 0.5 percentage points, on average, with better tax preparers offering larger reductions and charging higher fees. The size of tax savings in levels and as a share of gross income also increases in income taxpayer rank. We further find that tax preparers who are better at lowering their clients' tax obligation are also suspected more frequently of under-reporting by the tax authorities. Our results highlight the significant role that paid tax preparers play in shaping post-tax income disparities.

JEL codes: H20, H22, H23, H24, H26
Key words: Tax avoidance, tax preparers, inequality, tax revenue collection

[*]havivi@berkeley.edu, *Princeton University*

[†]kat.bilicka@usu.edu, *Utah State University, NBER, CEPR and Oxford University Centre for Business Taxation; Jon Huntsman Business School, Logan, UT, United States,*

[‡]jakob.brounstein@gmail.com, *Institute for Fiscal Studies*

[§]felipe.lobel@duke.edu, *Columbia University*

[¶]Alexander.Yuskavage@treasury.gov, *U.S. Department of the Treasury, Office of Tax Analysis.* This research was conducted while Yuskavage was an employee at the U.S. Department of the Treasury, and at no time was confidential taxpayer data ever outside of the Treasury or IRS computing environment. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily reflect the views or the official positions of the U.S. Department of the Treasury or the Internal Revenue Service. All results have been reviewed to ensure that no confidential information is disclosed. We thank seminar participants at the Institute for Fiscal Studies 2024 Research Away Day, the International Institute for Public Finance 2024 Annual Congress, Irish Public Economics Workshop, the and at ZEW 2024 Public Finance conference for valuable feedback. We acknowledge no sources of research support, financial relationships or other potential conflicts of interest that apply to any of the authors.

1 Introduction

Over half of American income tax filers use a paid tax preparer every year.[1] Tax preparers are specialized professionals who help their clientele (individual or business taxpayers) maintain tax compliance and minimize their tax burdens by combining their knowledge of the tax code (obligations, deductions, credits, legal structures, etc.) with clientele characteristics. In countries that feature automatic income tax filing, paid tax preparers tend to assist only taxpayers with more complicated income types; in countries that do not feature automatic income tax filing for most taxpayers, like the United States, paid preparers are widely used across the population.

However, despite widespread usage of tax preparers, the literature provides little evidence of their impact on tax revenues and on the distributional incidence of the income tax system.

In this paper, we study the role of tax preparers in tax optimization using the universe of US individual income tax returns between 2011 and 2020. First, we propose a novel measure of declaration-level tax optimization, which we use to measure the causal effects of tax preparers on individual income tax optimization. We exploit preparer-clientele switches to estimate preparer fixed effects and use those fixed effects to study the impacts of preparers on tax revenues and post-tax income inequality. Second, we explain *how* tax preparers engage in tax optimization for their clientele. To this end, we 1) analyze tax preparer characteristics that are correlated with tax preparer fixed effects and 2) study how the magnitude of tax optimization varies depending on whether tax preparers use deductions and/or credits to reduce their clients tax liability. Third, we study to what extent tax optimization and audit selection are related by estimating preparer-specific audit selection fixed effects and correlating the estimates with our preparer-level tax optimization fixed effects.

We begin by developing a novel measure of tax optimization that relies on the comparison between actual tax paid and the predicted tax paid based on a measure of gross income. We construct predicted taxes paid using a mapping from "broad" income and taxpayer charac-

[1] In fiscal year 2018, the US Internal Revenue Service's Statistics of Income program ("IRS SOI", "SOI") reported that 52.9% of income tax filers used a paid tax preparer. This figure specifically refers to *human* tax preparers as opposed to online tax preparation services, such as TurboTax or H&R Block.

teristics, such as state of residence, filing type, and year.[2] Our preferred tax optimization measure is constructed as the percentage point difference between the effective tax rate (ETR) predicted from broad income and final ETR (measured as the ratio of final taxes paid and broad income), which is then further residualized on income and year fixed effects. This measure has several advantages over previously used indicators of tax optimization: (1) unlike simple ETR, it allows for a within-return comparison of how much an individual is optimizing their taxes, rather than relative to other individuals or over time, (2) it is not mechanically related to income through the parametric shape of the income tax schedule based on income, and (3) when multiplied by gross income, it gives a prediction of levels of taxes "saved" relative to the predicted baseline.

To motivate the importance of tax preparation for US tax revenues, we present a set of new stylized, descriptive facts on the use of tax preparers across US taxpayers and their contribution to tax optimization. We find that consistently across all the sample years, around 65% of US gross fiscal income is handled by paid tax preparers. In 2020, this figure amounted to around $7.5 trillion (35% of US GDP). The use of tax preparers is prevalent across the entire income distribution, with around 55% of taxpayers using one, on average. Nevertheless, there is variation in tax preparer usage across the income distribution. Richer taxpayers are more likely to use tax preparers, with 90% of the top one percent of the richest taxpayers having one. Tax preparers' clients have lower ETRs conditional on their income, and their tax savings increase across the income distribution. Individuals at the 90th percentile have 2-3 pp lower ETRs if they use a tax preparer than if they do not. This descriptive "preparer ETR premium" sharply declines at the very top of the income distribution.

To quantify the causal contribution of tax preparers to tax optimization, we estimate a two-way fixed effect model with taxpayer and tax preparer fixed effects, inspired by a large literature in labor economics that models wage equations as an additive function of worker and firm effects (Kline et al., 2020; Card et al., 2013a; Abowd et al., 1999). The effects

[2] Our framework is analogous to the TAXSIM model of Feenberg and Coutts (1993); we measure broad income as the reported total income on the tax form and add the other non-taxable net income sources from the additional schedules C, E, F, and SE that include other kinds of business and miscellaneous income.

of taxpayers and tax preparers on tax optimization are identified among the connected set of taxpayers who switch between different tax preparers. For causal interpretation, identification assumption requires no systematic time- or individual- correlation between unobserved individual characteristics and preparer quality switches. Following the standard diagnostics proposed in the literature, we verify, using event studies, that payer-preparer match effects do not characterize taxpayers to preparers' assignment. Specifically, we find no pre-event trends in tax optimization around taxpayer switches to different tax preparers, and that switches to better or worse preparers are unaccompanied by systematic changes in income. Moving to a better tax preparer increases tax optimization, and moving to a worse preparer produces a symmetric reduction in tax optimization.

A one standard deviation better tax preparer reduces the ETRs of their clients by 0.5 percentage points. However, the magnitude of this tax saving varies across the income distribution, with tax preparers being most effective at the top, reducing ETRs by over 0.6 percentage points. We then show that taxpayers with highest income use the most effective tax preparers and that the top characteristic that explains by far the most variation in tax preparer effectiveness is how much they charge for their services; the other two important correlates are tax preparer age (with younger preparers optimizing more) and client churn (with those with less churn optimizing more). We proceed to examine how tax preparers save their clients money. We separately study the role of deductions and tax credits optimization to find that a 1 standard deviation better tax preparer deducts on average 2.5% more from their clients tax return and uses 1.2% more credits to lower their clients' tax liability. Tax preparers are more effective at reducing their clients tax liability through deductions for middle-income clients and most effective using credits for their high income clients.

We use our results to perform several benchmarking exercises that allow us to understand the role of tax preparers for tax revenues and income inequality. We find that eliminating the paid preparer industry completely would increase tax revenue by $54.4 billion and reduce the share of post-tax income held by top 1% by 0.08 percentage points. To contextualize this result, we observe that the income share of the top 1% of earners (measured by pre-tax broad income) only decreases by 1.15 percentage points. Therefore, this result implies that eliminating the paid preparer industry would augment the inequality-reducing properties of the US federal income tax system by around 7%. At the same time, assigning everyone in the top 1% of the income distribution a null-preparer would raise $20.5 billion in tax revenues and reduce the post-tax income share of the pre-tax top 1% of earners by 0.15 percentage points. Our results suggest that tax preparers play a large role in shaping post-tax income inequality in the US.

How expensive is it to use a paid tax preparer? From 2012 to 2017, taxpayers who itemized their tax deductions could also deduct fees paid for tax preparation services. While relatively few (around 20% of taxpayers) taxpayers itemized their deductions in a manner correlated with income, these data allow us to understand, descriptively, the costs of tax preparation, the tax savings that tax preparers generate net of their fees, and the correlation between fees and preparer fixed effects. An average tax preparer fee in our data is USD 400. Fees appear constant income at around USD 200 until the 80th percentile of the income distribution. At this point, average fees increase sharply; we observe an average fees within the top 1% of earners of about USD 3000. The net tax saving generated by tax preparers increases across the income distribution monotonically, varying from nearly USD 50 at the bottom to almost USD 8000 for the top 1 percentile of the income distribution. By correlating fees with tax preparer fixed effects, we find that better tax preparers charge larger fees and generate larger net tax savings. We also find that nearly all of the variation in preparer fees are explained by variation *across* preparers as opposed to variation in the services provided by a fixed preparer.

We conclude, by studying preparer-level audit selection as a measure of suspected evasion. We estimate our baseline two-way fixed effects model using audit selection as a dependent variable. Since tax authorities target audits to result in detection of evasion, we interpret the dependent variable as indicating a tax return that is highly suspected to belong to a tax evader.[3] We find that the preparer-specific audit fixed effects are positively correlated with preparer-specific tax optimization fixed effects. This fiding suggest that tax optimization is linked with tax evasion. Because tax preparers are not assigned randomly, we cannot causally

[3] Indeed, 85% of operational audits result in an adjustment.

separate between the possibilities either that tax-evading taxpayers sort into high-optimizing preparers, or that high-optimizing preparers generate tax evasion.

1.1 Contributions to the Literature

Our paper contributes to the literature analyzing the role of paid tax preparers for tax compliance both for individuals and firms. The paper closest to our work is Breunig et al. (2024) who examine the role of paid tax preparers in bunching behavior of individual taxpayers in Australia. The authors show that tax preparers are more likely to deliver positive round number bunching refunds and the main mechanism through which they lower their taxpayers income is deductions. Our paper differs from theirs in two dimensions. First, we estimate preparer fixed effects, rather than simply control for them, to understand the role of preparer quality for tax optimization. Second, we focus on tax optimization across the entire income distribution, not just for those individuals who bunch.

Most of the remaining existing work in this space considers the mediating role of tax preparers during audits or in take-up of tax credits or tax refunds (Battaglini et al., 2022; DeBacker et al., 2024; Boning et al., 2020; Zwick, 2021; Kopczuk and Pop-Eleches, 2007; Goldin et al., 2022). Specifically, in the US (Boning et al., 2020) and in Italy (Battaglini et al., 2022) show the role of tax preparers as information hubs through which behavioral effects of tax audits affect other taxpayers in the same network. In the US, Kopczuk and Pop-Eleches (2007); Goldin et al. (2022) discuss the role of tax preparers in facilitating the take-up of Earned Income Tax Credit or Child Tax Credit. Also in the US, Zwick (2021) uses similar research design to ours to show that better quality tax preparers increase claim refunds for tax losses for small firms.[4] These papers examine the role of tax preparers for the behavior of small firms, while we focus on individual taxpayers. DeBacker et al. (2024) show that a higher adjustment amount is paid by individual tax preparer users following an audit. While we also consider the effects of audits in this paper, we use these audits in an auxiliary analysis to disentangle the tax optimization from tax evasion.

[4] A related paper by Belnap et al. (2024) examines the role of individual tax preparers for corporate firm outcomes to show that individual tax prepares matter more than firms that employ them, but that internal actors, such as top accountants, and executives are more important for corporate tax planning.

This paper also contributes to the large literature describing inequality, tax strategy (both evasive and avoiding) and the distributional incidence of the US federal income tax (Auten and Splinter (ming); Guyton et al. (2023); Smith et al. (2022); Piketty et al. (2017); Saez and Zucman (2016); Piketty and Saez (2003)). These paper have sought to characterize the evolution in income and wealth inequality over time as well as the incidence of tax noncompliance issues across the income distribution (also Brounstein (2023); Londoño-Vélez and Ávila Mahecha (2021); Alstadsæter et al. (2019); Alstadsæter et al. (2018)). Beyond characterizing the value added of tax preparers, our work informs our understanding of income inequality and the US tax code by describing the distribution of tax preparer value added across the income distribution. How does the institution of paid tax preparation service cause the *de facto* progressivity of the US federal income tax code to depart from its statutory progressivity? Our work also seeks to estimate the counterfactual scenarios of income inequality under different allocations of tax preparer value added. Finally, we will identify the exact provisions that facilitate both legal and illicit tax optimization.

2 Measuring tax Optimization

Tax preparers are specialized professionals who operate for the purpose of helping their clientele (individual or business taxpayers) maintain tax compliance and minimize their tax burdens. Tax preparers can come from a variety of career backgrounds—e.g. from law, from accounting, etc. They combine their knowledge of the tax code (obligations, deductions, credits, legal structures, etc.) with clientele characteristics in order to reduce the amount of tax that their clients have to pay. In principle, tax preparers help their clientele with all kinds of taxes, but are understood to most commonly focus on taxes on income, which, in high-income countries, tend to be the highest and most salient to taxpayers.[5]

In countries that feature automatic income tax filing, paid tax preparers tend to primarily assist taxpayers with more complicated income types. Such taxpayers tend to have higher

[5]The US Census Bureau finds that the tax preparation industry (identified as 2022 NAICS code 541213) generated USD 7.25 billion in revenue in 2022 (Bureau, 2022). IBISWorld, a private industry statistics company, gives an estimate of 2024 revenue of the tax preparation industry of USD 13 billion, although their methodology is proprietary and it is not clear how it may differ from that of the US Census Bureau.

incomes that arise from a variety of income sources. In countries that do not feature automatic income tax filing for most taxpayers, like the United States, paid preparers are widely used across the population. However, most high income countries—unlike the US—feature a system of simplified, automatic personal income tax filing on behalf of most taxpayers with relatively simple income portfolios (e.g. entirely third-party-reported wage earnings). While countries with automatic personal income tax filing likely see lower usage of tax preparers across the income distribution, such countries are also more likely to see much more concentrated usage of preparers among the highest earners, due to their more complicated income composition. To the extent that we find that tax preparers generate largest tax savings at the top of the income distribution, this means that our findings are valid outside of the US and can be used to understand the effects of tax preparers on top-earner income inequality.

2.1 Key concepts: Tax preparers and broad income

Tax preparers. A key feature of our data is that taxpayers utilizing professional tax services report their preparer's tax identification number on Form 1040. Since 2011, paid tax professionals are required to register with the US Internal Revenue Service (IRS) and obtain a fixed ID number. This feature enables us to track each individual's history of preparers usage including how often they change tax preparers. Further, because we can identify who the tax preparers are, we are also able to observe their clientele over time as well as analyze the tax preparers themselves.

Importantly, these tax preparer ID numbers correspond with *individuals who are tax preparers*, not with the firms that employ such individuals.[6] Further, these ID numbers are unpopulated when individuals use either paid or unpaid tax preparation and filing software that is unassisted by an actual tax preparer, such as TurboTax or H&R Block.[7]

We index all tax preparers as $j \in \mathcal{J}$ given their observed ID. We hard-code all individuals

[6] We *can* assign tax preparers to de-identified firms by observing whether a tax preparer has a third-party-reported wage income declaration (form W-2), and tagging the de-identified ID number of the employing firm.

[7] We can identify the usage of such software in some cases. Prior to 2018, individuals itemizing their deductions could deduct tax preparation fees; in the case that we observe such fees without a paid tax preparer present, we categorize these individuals as using tax preparation software.

without a paid tax preparer to the tax preparer ID code of $j = 0$, which we term "the null-preparer", i.e. self-preparation. This null-preparer class contains all taxpayers who do not use a paid a tax preparer, whether they use a tax preparation software or manually self-prepare and self-file. Additionally, we also hard-map all Volunteer Income Tax Assistance (VITA) or Tax Counseling for the Elderly (TCE) preparer users to a separate preparer ID.[8] Indexing all taxpayers as $i \in \mathcal{I}$ in year $t \in \mathcal{T}$, we can define a function $J(i, t)$ that maps taxpayer i and year t to a single tax preparer j. Importantly, the mapping of taxpayers to tax preparers is a well-defined function, where only a single preparer at most can be listed on an taxpayer's income tax return.

Broad income. For all taxpayers, we assign a gross measure of earnings we call "broad income" that aggregates earnings across wages, capital income, farm income, self-employment income, partnership income, and other business income (e.g. from S-corporations.).[9] This measure is broader in scope than the pre-deduction measure of total income used by the IRS and allows us to measure gross income prior to the application of credits and deductions and organization of income into different earnings vehicles. This has two advantages: (1) it abstracts from the earnings source, which can see differential tax preference (e.g. long term capital gains) and (2) it abstracts from tax strategies that shift *between* income sources. In principle, this income concept, which we denote B, represents earnings exogenous from tax strategy.[10]

2.2 Measuring tax optimization.

Tax optimization is the process of lowering one's tax obligation, either through evasion or by using the tax code provisions to engage in *legal* tax avoidance. There does not exist

[8] These programs connect low income and elderly (for TCE) taxpayers with volunteer tax preparers. The tax preparers themselves are not identified in form 1040 where a paid preparer would place her IRS ID number, but we observe the usage of these programs.

[9] We construct this income concept using revenues, expenses, and realized capital gains/losses from the form 1040 main form as well as its schedules SE, C, D, E, and F, as well as negative amounts reported on additional schedules.

[10] There may be endogenous real activity responses to anticipated tax optimization strategy, as that strategy affects the actual income tax rate that individuals pay. We abstract from such endogenous responses, as evidence demonstrates small earnings elasticities to the net-of-tax wage rate (Kleven et al., 2024; Tortarolo et al., 2020) and even smaller responses to low-saliency shocks or inattention (Boccanfuso and Ferey, 2023).

a consensus around a quantitative measure of tax optimization. Typically, when making interpersonal comparisons or intrapersonal comparisons over time, researchers use effective tax rate (ETR), usually defined as final taxes paid divided by a gross measure of income. However, there are at least two reasons why ETR may not be suitable for quantifying tax optimization in our context. First, ETR is a nearly deterministic function of income and we want to evaluate tax optimization independent of income. Second, ETR is uninformative of tax optimization in cross-sectional settings as it offers no within-individual baseline to compare against. For example, an ETR of 12% does not inform the researcher how much a taxpayer is "tax optimizing" unless one compares to a similar taxpayer or over time with reference to an ETR change.

Instead, in this paper we propose a measure of tax optimization based on comparison between imputed taxes at the "start" of the income tax filing process and the final taxes paid at the "end" of the income tax payment process. Specifically, we design a mapping $T^b = T(B;\theta)$ from broad income B to federal income tax liability T that considers broad income B as ordinary wage income and takes filing characteristics as state of residence, filing type, and year as a vector of covariates $\theta \in \Theta$, following a simplification of TAXSIM (Feenberg and Coutts, 1993).[11] This mapping features the application of deductions, D, that remove income from tax base and credits, C, that remove tax from calculated taxes to arrive at a final tax liability, which can be defined as $T^f = T(B - D;\theta) - C$. We define a non-counterfactual benchmark "tax savings" as $T^b - T^f$. Note that this measure ultimately corresponds with the difference between imputed statutory tax obligation using our broad measure of income and actual taxes paid.

Using our mappings, we define our main measure of tax optimization as:

$$\text{ETR difference} := \frac{T^b - T^f}{B} = \frac{T^b}{B} - \frac{T^f}{B}. \tag{1}$$

ETR difference has several tractable and useful interpretations. First, as presented in Equation (1), ETR difference corresponds to tax savings as a share of broad income. This

[11] By keeping characteristics such as filing status and state as fixed, we implicitly shut down tax optimization on these margins.

definition means that for a given tax return multiplying a value of ETR difference by a return's broad income yields a measure of tax savings. Second, this measure also corresponds with the percentage point (levels) difference in implied ETR using tax imputed statutorily from broad income and actual ETR. E.g. an ETR difference value of 0.02 means that one's final ETR was 2 percentage points lower than would have been based on the statutory income tax schedule applied directly to broad income, keeping other filing characteristics fixed.[12]

Note that this measure does not quantify tax optimization strategy relative to a true counterfactual. Instead, it is useful in establishing accounting benchmarks and for setting up more tractable comparisons between taxpayers. This measure likely still demonstrates *some* correlation with income insofar as tax optimization may have different returns at different parts of the income distribution, which also has different ETR. However, this measure is not mechanically related to income to the same extent as is ETR.

3 Data and descriptive statistics

We use data on the universe of US annual personal income tax returns, which allow us to track the income of individual taxpayers and their tax preparers over time. We construct our dataset using individual income tax Forms 1040 filed by US individuals and taxpayers annually between 2011 and 2020. This rich dataset includes detailed information on income, deductions, and credits, along with supplementary schedules that break down these amounts into more specific items.

We apply two sample restrictions on the full population of the US individual taxpayers. First, we remove taxpayers who earn below the standard income tax deduction. These individuals mechanically optimize 100% of their income through the standard deduction. Moreover, below this threshold, our tax optimization measures given their definitions, are only well-behaved within the set of income tax obligations above zero.[13]

[12] We consider alternative measures of tax optimization and replicate our main results with those measures in Section C.

[13] Considering the effects of tax preparers on tax optimization for Earned Income Tax Credit (EITC) filers would be a great extension of Kopczuk and Pop-Eleches (2007) and is a promising avenue for future research.

The second sample restriction drops individuals who are excluded from the leave-one-out largest connected set of taxpayers and tax preparers. Our central methodology in Section 4 involves identifying taxpayer and tax preparer optimization fixed effects using a series of two-way fixed effects regressions that are identified on the largest sample of taxpayers and preparers that are "connected" by switches between the two groups—the "largest connected set" (Abowd et al., 1999). We implement the methodological refinement from (Kline et al., 2020) that uses a more restrictive sample of taxpayers, the "leave-one-out largest connected set", which consists of the largest connected set such that every tax preparer remains connected after removing any single taxpayer from the sample.[14]

Table 1 summarizes baseline descriptive statistics for our sample, showing characteristics of taxpayers in panel A (splitting by tax preparer users and non-users) and tax preparers in Panel B. Our final sample includes over 175 million unique taxpayers and 1.268 million unique tax preparers, yielding almost 1 billion taxpayers-years from 2011 to 2020. Note that we observe different taxpayer types, such as individuals or spouses, which we do not disaggregate into individual taxpayer units (e.g. we do not split earnings by joint-filers 50-50 to both spouses). On average, taxpayers who used tax preparers in this period had an annual gross income of about \$118,000 and broad income of \$158,000 (2019).[15] Taxpayers who did not use tax preparers have much lower annual gross incomes of \$75,000 and broad income of \$91,000. Taxpayers who use tax preparers are on average older relative to those who do not use tax preparers, and face an average ETR of 7.8%, which is 8.7 percentage points lower than the average ETR otherwise statutorily paid on broad income . Non-tax preparer users face ETRs that are only 7.4 percentage points lower than the counterfactual. Panel B shows that paid preparers are predominantly male, with average annual income of over \$96,000. Typically, these preparers serve 63 clients annually and there is a 37% chance of retaining a client from one year to the next.

[14]Table A.1 gives basic descriptive statistics about this sample restriction, showing that around 25% of taxpayers and 22% of unique tax preparers are dropped by imposing this requirement which represented 15% of taxpayer-year observations and 0.05% of taxpayer-year observations respectively. Excluded taxpayers earn USD 30,000 less per year than included taxpayers. Panel (b) gives basic statistics about preparers, showing that excluded preparers have on average only 1 client per year, whereas those included have about 77 clients per year.

[15]We express all monetary values, unless otherwise explicitly specified, in real USD 2019.

Table B.1 shows that there is substantial taxpayer-level variation in the client-preparer match. Columns (3) and (4) demonstrate that a large share of clients use multiple different preparers over the sample period, and 50% of them use at least one different professional paid preparer. Table B.3 displays descriptive statistics on preparer rotation among clients. On average, clients stay with a preparer for around 3 years. In any given year, we observe a tax preparer switch probability of around 20%, where clients either switch from preparer-to-null, null-to-preparer, or preparer-to-different-preparer. 54% of switches are changing between two paid preparers. Figure B.1 shows that the probability of preparer switches is similar (and converging) between income quintiles, with by 2020, the top fifth of earners exhibiting around a 14% switch probability and the bottom fifth of earners displaying an 18-19% switch probability.[16]

3.1 Stylized Facts

We begin our analysis by demonstrating four novel facts about preparer usage and tax optimization throughout the income distribution.

Fact 1: Tax preparers usage is widespread. During our sample period, around 55% of the US taxpayers have used a tax preparer and about 65% of the US gross fiscal income has been handled by a paid tax preparer. In fact, tax preparers handled about \$6 trillion of aggregate gross income (AGI) in 2011 and about \$7.5 trillion in 2020. Around 550,000-600,000 tax preparers handle individual income tax returns each year of the sample (Figure A.1 that demonstrates the stability of these figures across years).

Fact 2: Richer taxpayers are more likely to use a tax preparer. In Figure 1 Panel (a) we document how preparer usage varies throughout the distribution of broad income. The probability of preparer usage increases approximately monotonically in income with the median taxpayer having 50% probability of using a paid tax preparer. This probability

[16] We do not observe any notable changes in the likelihood of tax preparer switches around the TCJA reform that changed the deductibility of tax preparer fees. We discuss the fee datasets and results in detail in Section 5.

increases slowly at first, to 60% at the 90th percentile of the income distribution, displaying a sharp increase between the 90th and the 99th percentile of income distribution from 60% to 90% of individuals reporting to use a paid tax preparer.[17]

Fact 3: Taxpayers who use tax preparers have lower ETRs. Figure 1 Panel (b) shows that preparer users face a lower ETR over most of the income distribution. As expected, ETR increases mechanically in income rank. At the bottom of the income distribution, up until the thirtieth percentile, there does not appear to be a large difference in ETRs between paid preparer users and non-users. At this point, a small difference of around 1 percentage point emerges, with non-users paying higher ETRs. This descriptive preparer ETR premium begins to widen at the 80th percentile of the income distribution to around 2 percentage points around p97, and then sharply declines at the very top of the income distribution. Panel (c) estimates this within percentile difference explicitly and shows that the descriptive preparer ETR premium is entirely eliminated at the very top of the income distribution. This figure highlights the likely selection into or out of having a paid preparer at the very top of the income distribution, perhaps either because those individuals are preparers themselves or receive tax preparation services through other means.

Fact 4: The size of the tax preparer network varies across the income distribution. In Panel (d), we plot the average log count of total "peer clientele" shared by tax preparer user, i, which is the count of other individuals $k \in \mathcal{I} \neq i$ such that $J(k,t) = J(i,t)$. The figure reveals an inverted right-skewed U shape, where tax preparer peer networks increase within the interior of the income distribution until around the 80th percentile. Tax preparers who serve clients at the thirtieth percentile of the income distribution have on average 180 clients, those serving clients at the 80th percentile have 250 clients, but those serving clients at the top percentile of income distribution have on average 160 clients with the number dropping for those at the top 0.01 percentile to below 90 clients (Figure A.3 Panel (d)). This pattern further highlights potential selection at the top of the income distribution

[17]Figure A.2 disaggregates this figure based on whether non-wage income exceeds half of broad income. Non-wage earners exhibit an approximately 20 percentage point higher probability of using a page preparer than wage earners; this difference appears stable across the income distribution.

and could be consistent with the specialized nature of services for high-income clients that limits the scope for expanding operations. Conversely, preparers serving lower-income clients often operate smaller businesses with limited capacity for scaling up their client base.[18]

3.2 Tax optimization descriptive statistics

In Figure 2 we describe the main tax optimization measure that we use in detail. First, in Panel (a), we plot the full distribution of ETR difference for all taxpayer-years in our sample. ETR difference ranges between 0.03 and 0.16, with maximum density at approximately 0.05, decreasing monotonically from that point until 0.16. Note that we observe some bunching at 0.10, reflecting an outsized mass of taxpayers that face a statutory ETR of 10% who entirely remove their tax obligations.

Panel (b) illustrates the distribution of ETR difference across broad income rank. Tax optimization, as measured by ETR difference, is largely flat from the 20th percentile to the 90th percentile of the income distribution at around 7 percentage points. This means that taxpayers in this income range face an approximately 7 percentage points lower ETR than the one imputed statutorily based on their broad income. At the bottom of the income distribution we see a higher level of ETR difference—between 8 and 10 percentage points—as well as substantial variability.[19] At the top end of income distribution, the ETR difference begins to increase sharply starting at around the 90th percentile of the broad income distribution. The top 1% of taxpayers have an average ETR difference of 17 percentage points.

Panel (c) corresponds with Panel (b), separating ETR differences based on paid preparer usage. We find no descriptive tax optimization premium for the bottom 20 percent of the income distribution, but at this point, the premium widens to around 1 percentage point, where the premium remains constant between the 40th and 80th percentiles of the income

[18] Figure A.3 provides additional information on preparer clientele. Panels (a) and (b) document small bunching at round numbers, suggesting that some preparers have a fixed number of clients that they are willing to accept in a given year. Panel (c) shows the increasing sizes of tax preparer networks over time.

[19] This greater variation likely occurs due to smaller sample size, as individuals at the bottom of the income distribution are disproportionately dropped by our minimum income requirements and inclusion in the leave-out-out largest connected set. However, note that the bottom of the income distribution also benefits from means-tested deductions and refundable tax credits, such as the EITC.

distribution. The premium widens within the top fifth of the income distribution to around 2 percentage points before diminishing at the very top. Note that this pattern is very similar to what Panel (b) of Figure 1 indicates when using ETRs. Panel (d) breaks out the top 1% of earners, and finds that the preparer premium actually becomes negative within the top 0.5% of earners. For the top 0.01% of earners, the gap widens considerably in favor of taxpayers *without* a paid preparer, to around 7 percentage points and is consistent with the de facto regressivity of the US federal income tax system for the very top earners (Piketty et al., 2017).

Note that our main measure of tax optimization displays small unexpected local non-monotonicities (e.g. the small inverse U-shape in ETR difference in panels (a) and (b) between p45 and p75). These non-monotonicities likely reflect local interactions between ETR and income and changes in tax optimization opportunities that differ along the income distribution (e.g. the mortgage interest deduction which depends on home ownership). This characteristic along with the reversal of the descriptive preparer tax optimization and ETR premium at the very top of the income distribution highlight the descriptive nature of these relationships. Consequently, in the next section, we investigate the role of tax preparers using a causal design, which allows us to separate individual effects and other confounders from tax preparer effects.

4 The causal effect of tax preparers

To identify the causal contribution of tax preparers to tax optimization, inspired by Abowd et al. (1999) and Card et al. (2013a), we exploit moves of taxpayers between different tax preparers over time. Our analysis includes both moves between different preparers and taxpayers moving to and from having no preparer to using the services of a new tax preparer.

Formally, we have a population of N taxpayers, indexed by i, who in every year $t \in \{1, .., T\}$ are characterized by tax optimization level, z_{it}, and other individual, possibly time-variant, characteristics, x_{it}. Here, we use the tax optimization measure described in Section 3. At the end of each calendar year, taxpayers decide whether to file their taxes alone or use the services of one of $\mathcal{J}$ tax preparers indexed by j, where $j = 0$ indicates using no tax preparer (the null-preparer). This last condition allows us to estimate the effectiveness of tax preparers relative to the null-preparer alternative. Importantly, the null-preparer alternative includes all taxpayers who do not use a paid tax preparer, including, for example, taxpayers who use free or paid online tax filing software or those who file manually for free. We use the function $J(i,t) \in \{0, 1, ...\mathcal{J}\}$ to describe the tax preparer chosen by taxpayer i in year t.

We model the tax optimization z_{it} of taxpayer i at time t as an additive linear function of tax preparer component $\psi_{J(i,t)}$, taxpayer component α_i, and an index of observable time-variant characteristics $x'_{it}\beta$:

$$z_{it} = \alpha_i + \psi_{J(i,t)} + x'_{it}\beta + r_{it}, \tag{2}$$

where r_{it} is the error term, ψ_j is preparer j's effectiveness that is common to all their clients, and α_i is taxpayer i's skills, productivity, and time-invariant characteristics that affect both their earnings and their tax saving. Lastly, x_{it} includes the logarithm of broad income, year fixed-effects, and the taxpayer's state of residence that varies by year, which accounts for changes in the state-specific tax laws and any resulting tax preparation specialization changes.

We estimate Equation (2) by Ordinary Least Squares (OLS). Crucially, the tax preparer and tax payer parameters ψ_j and α_i are identified only within a connected set of tax preparers that are linked by taxpayers' mobility, i.e., those that move at least once between the tax preparers during the sample period (Abowd et al., 1999).

Identification of preparer and taxpayer fixed effects from Equation 2 are premised under the assumption that the assignment of taxpayers to preparers is uncorrelated with unobserved payer-preparer matching. Formally:

Assumption A1. *(Strict Exogeneity)*

$$E[r_{it}|X, \psi, \alpha] = 0,$$

where $X = (x_{11}, ..., x_{NT})$, $\psi = (\psi_0, .., \psi_{\mathcal{J}})$, and $\alpha = (\alpha_1, .., \alpha_N)$. This assumption allows for unrestricted dependence of preparer mobility on tax preparers and taxpayers' effects. For example, highly skilled taxpayers may be more likely to use the services of highly effective tax preparers than low-skilled taxpayers. However, assuming mean independence between the error term and the entire path of future preparers rules out mobility with respect to these time-varying and/or match effect shocks.

4.1 Targeting the variance

Estimation of Equation 2 results in two vectors: $\hat{\psi} = (\hat{\psi}_1, ..., \hat{\psi}_J)'$, and $\hat{\alpha} = (\hat{\alpha}_1, ..., \hat{\alpha}_N)$, which are, under Assumption A1, unbiased estimates for $\psi = (\psi_1, ..., \psi_J)'$ and $\alpha = (\alpha_1, ..., \alpha_N)$, respectively. Our main target parameters are the square root of the variance component of ψ and α, i.e., the standard deviations of individual and preparer effects in our sample. To quantify the role of sorting, we additionally report the correlation between tax preparer and taxpayer effects, which describes the extent to which more sophisticated taxpayers are more likely to sort into more effective tax preparers.

Each estimated effect, $\hat{\psi}_j$, is an unbiased, albeit noisy, estimate of the true preparer effect. As such, estimating the variance component with its empirical analog, $\sum_j \frac{\hat{\psi}_j - \bar{\psi}}{J-1}$, produces an upwards-biased estimate for the variance components. As noted by Kline et al. (2020) and Bonhomme et al. (2023), such bias could be non-negligible. To correct for this bias, we apply and report the Kline et al. (2020) (hereafter, KSS) leave-out unbiased estimate. The KSS estimator is identified only within a leave-one-out connected set sample, which is the sample in which the set of taxpayers and tax preparers remain connected even after each taxpayer (i, t) has been taken out.[20]

4.2 Validating identification assumptions

If the variation across tax preparers is mainly attributable to sorting on unobserved characteristics, then taxpayers who switch between tax preparers will not necessarily experience a

[20]Kline et al. (2020) provides the conditions under which the leave-one-out variance component is a consistent estimate for the true variance component.

proportional change in their tax optimization measures. If, on the other hand, tax preparers are characterized by a stable latent effectiveness applied to all of their clients, then we might find that taxpayers who switch to a tax preparer whose other taxpayers experience higher tax optimization gains, will experience a gain as well, regardless of their original preparer.

Following this intuition, Card et al. (2013a) propose a simple empirical test to evaluate the plausibility of this assumption. Figure 3 plots a series of simple event studies that analyze taxpayers' transitions to different tax preparers. We assign tax preparers to their quartile of the distribution of clientele leave-out mean optimization,[21] where 1 indicates bottom 25th percentile in the distribution of the tax optimization (least effective) and 4 indicates top 25th percentile (most effective).

We then track clients' tax optimization over time, residualized on the other covariates in Equation (2), both before and after they switch preparers. We find that clients transitioning from the least to most effective preparers experience significant increases in their tax optimization. Conversely, moving from most to least effective tax preparer results in a symmetric reduction in tax optimization. The figures also suggest that different types of tax preparer movers experience different levels of tax optimization benefits and losses. Furthermore, prior to switching preparers, clients demonstrate a stable tax optimization behavior, suggesting no differential pre trends.[22] These findings validate the exogeneity assumption.[23]

4.3 Estimates of variance components

Panel (i) of Table 2 shows that tax preparers and taxpayers contribute substantial variability to tax optimization. A one standard deviation increase in tax preparer quality yields

[21]I.e. for a taxpayer k with preparer j and a set of $K-1$ other taxpayers $-k$ that also employ tax preparer j in year $\bar{t}$, the leave-out mean optimization of taxpayer i's preparer j is defined as $\frac{1}{K-1}\sum_{i\in -k} x_{i\bar{t}}$.

[22]Figure B.2 displays a version of this figure using the raw, non-residualized version of ETR difference as the outcome and preparer-ranking variable. The figure shows similar results; however, all units regardless of their initial or final preparer tax optimization rank, appear to be on a common slight downward trend around their move. This common trend does not threaten our identification. We interpret this downward trend as reflecting a drop in aggregate tax optimization following TCJA. Because TCJA occurs later in our panel, this drop in aggregate tax optimization would manifest as a common downward trend.

[23]Figure B.3 shows a placebo test for further testing our strict exogeneity assumption, plotting the evolution in taxpayer broad income around moves to different preparer effectiveness ranks. The figure shows that taxpayer switches to different preparer quality broadly are not characterized by differential effects on income.

a half percentage point increase in ETR difference. The contribution of tax payers is even larger; a one standard deviation increase in tax payer fixed effect yields a 2.7 percentage point increase in ETR difference. Interestingly, Table 2 reports a very small positive correlation between the fixed effects of taxpayers and tax preparers. Unlike the literature in labor economics on worker-firm sorting on wage effects, which documents strong sorting, our results indicate that there is very little sorting of sophisticated or high-optimizing taxpayers with high-quality tax preparers. Column (2), which reports the same variance components for the sample of taxpayers and tax preparers excluding the null-preparer, shows similar qualitative estimates.

The contribution of the taxpayers and tax preparers to tax optimization increases slightly with taxpayers' income. Columns (3)-(5) display three sets of analogous estimates from three separate stratifications of our sample by within-year broad income tercile. The standard deviations of payers and preparers increase moderately with income terciles for both taxpayers and tax preparers. In contrast, we find no equivalent change in the correlation between the preparer and taxpayer effects by income.

4.4 Describing tax preparers

Figure 4 plots the distribution of tax preparer fixed effects across the income rank. To avoid the mechanical correlation between estimated fixed effect and taxpayer's income, we split the sample randomly into two samples, and for each taxpayer, we match the preparer fixed effect from the other sample. Each percentile rank includes approximately 10 million taxpayer-years or 1 million taxpayers per year, and therefore reports a parameter that is less noisy than each $\hat{\psi}_j$ itself. This figure suggests that the very high income taxpayers use the most effective tax preparers. In fact, taxpayers at the 20th percentile of the income distribution tend to use tax preparers who reduce their clients ETRs by approximately 0.15 percentage points. In turn, taxpayers at the 99th percentile of the income distribution tend to use tax preparers who reduce their clients ETRs by approximately 0.65 percentage points.

To further describe taxpayers, Table 3 reports the relationship between tax preparer effects and preparer characteristics, estimated via OLS regression. The most predictive characteristic of preparer effectiveness is the fee they charge; effective tax preparers charge higher fees. The other two important characteristics are the preparer year of birth and the probability of their clients staying next year, i.e., the client churn. We find that younger preparers and those with more stable client base optimize more. Taken together with evidence from Figure 4, these results suggest that tax preparers contribute to income inequality, as more affluent individuals are more likely to be able to afford to hire more effective tax preparers.

4.5 Margins of adjustment

Next, we turn to understanding how tax preparers affect different margins of tax optimization. We construct two additional measures of tax optimization: deductions and credits optimization. The US income tax form 1040 broadly features base/deductions optimization for taxable income[24] Following the calculation of initial tax liability in the individual income tax form, tax optimization then focuses on the application of nonrefundable and refundable credits. Consequently, we can directly infer both of those optimization types from tax returns without any imputations. Since the tax liability is calculated by applying deductions to the tax base first, we start with deductions optimization, which we define as:

$$Deduction\ optimzation = \frac{T(B;\theta) - T(B-D;\theta)}{T^b} = 1 - \frac{T(B-D;\theta)}{T^b} \tag{3}$$

After calculating preliminary tax liability based on deductions, tax credits can be used to further lower tax liability. We define credit optimization as follows:

$$Credit\ optimization = \frac{T(B-D;\theta) - (T(B-D;\theta) - C)}{T(B;\theta)} = \frac{T(B-D;\theta) - T^f}{T^b} \tag{4}$$

We use these two subcategories of tax optimization to consider *how* tax preparers optimize their taxpayers' ETRs. These measures indicate the proportional reduction in tax relative to tax statutorily imputed from broad income due to deductions and credits, respectively.[25]

[24] In the 2021 version of Form 1040, Taxable Income is defined on line 15.

[25] These measures also feature the property that their sum yields our measure of proportion tax optimiza-

We present evidence for deductions optimization in Panel (ii) and credit optimization in panel (iii) of Table 2. We find that a 1 standard deviation better tax preparer deducts, on average, 2.6% more from their client's tax return and they use 1.2% more credits to lower their taxpayers liability. This finding suggests that, on average, deductions are a more important mechanism through which the tax preparer can be more effective. This finding is in line with evidence from Breunig et al. (2024) from Australia who also highlight the role of deductions in their bunching results.

Further, our split sample results suggest that the mechanisms through which tax preparers optimize their client's tax returns vary across the income distribution. Tax preparers are most effective at reducing their client's tax liability through deductions for middle-income clients, while they are most effective at using credits for their high-income clients.

4.6 Inequality and revenue consequences

We now use our estimates to understand the inequality and revenue consequences of tax preparer usage across the income distribution. Using Equation (2), we perform simple counterfactual exercises that change the preparer fixed effect while holding all other factors fixed. To measure the economic impact the current assortative allocation of tax preparers relative to different benchmark scenarios, we construct a measure of economic impact that is related to, but distinct from, changes in collection of tax revenue. We compute our measure by multiplying the ETR difference (z_{it}) by broad income to obtain a measure of gross tax savings. This measure combines the overall breadth of income affected by the counterfactual with the degree to which that income makes a contribution to revenue collections. However, because this measure does not take into account the degree to which tax optimization is related to tax non-compliance or other misreporting we cannot say to what degree these impacts might translate into ultimate revenue collections

As defined above, gross tax saving is the difference between taxes imputed statutorily from broad income and final taxes paid: $T(B;\theta) - T^f$.[26] We then compute gross tax savings

tion developed in Section C.

[26] This measure does not consider fees, which we observe only for a selected sample of itemizing individuals.

of a tax preparer j relative to, for example, the null-preparer for an individual i in year t to be:

$$\text{Expected gross preparer tax savings}_{it} = \mathbb{E}[B_{it} \cdot (z_{it,J(it)} - z_{it,j=0})] = \mathbb{E}[B_{it} \cdot \psi_j], \qquad (5)$$

where expectations are taken over the prior distribution of ψ_j. We approximate this object by collapsing the data by income percentile ranks, therefore alleviating concerns about measurement error. Using this equation, we estimate different counterfactual scenarios of different allocations of preparers for different taxpayer groups.[27] Such counterfactuals represent purely mechanical benchmarks and do not take into account general equilibrium effects.

In Figure 5, we show the tax savings impacts of these counterfactual experiments, where we assign different tax preparers to either all taxpayers (in gray), all non-users of tax preparers (in green) and only the top 1% of taxpayers (in orange). We find that the decrease in gross tax savings from assigning all taxpayers a null tax preparer is around \$54 billions. This experiment is equivalent to mechanically removing the tax preparer industry altogether. Assigning a null tax preparer to only the top 1% of taxpayers still decreases their gross tax savings by \$20.4 billion due to the fact that higher-income taxpayers are both more likely to use a tax preparer and their tax preparers are more effective. Assigning all *non-users* a mean-ability tax preparer increases gross tax savings by \$19.7 billion, assigning all *taxpayers* a mean tax preparer increases gross tax savings by \$8 billion, and assigning the *top 1%* of taxpayers a mean tax preparer would generate a decrease in gross tax savings. This pattern occurs again because higher income taxpayers have more effective tax preparers who are more likely to save them more money than an average tax preparer.

In Table 4, we evaluate the consequences of different tax preparer assignments on post-tax inequality. As a baseline, we focus on top 1% share of post-tax income.[28] Our pre-tax

We discuss these issues and present results on fees in Section 5.

[27] Our alternate measures of tax optimization allow for similar expressions of gross tax savings, which we explore in the Appendix.

[28] Note that this exercise is slightly different from Piketty et al. (2017) who compare the top 1% income shares among the top 1% of taxpayers by pre-tax income and the top 1% of taxpayers by post-tax income separately. They also measure taxes other than federal income tax that are not relevant for tax preparer usage.

measures of income are similar in concept and magnitude: we measure the top 1% pre-tax broad income share at 24.2%, whereas Piketty et al. (2017) estimate this figure at 20.2%. We attribute the difference to the differences in income concepts employed and our sample restriction that disproportionately excludes lower earners. Our post-tax 1% income shares, however, are conceptually different and, in principle, measure the concentration of income in different groups. We find only a 1.15 percentage point reduction in the broad income share of the top 1% of pre-tax broad income earners following the payment of federal taxes.

Our results suggest that different hypothetical reallocations of tax preparers with different effectiveness to different taxpayer groups yield changes to the top 1% income share that are small in magnitude, but not necessarily small compared to this baseline. Assigning all taxpayers in the top 1% of pre-tax income to the null-preparer wold reduce their post-tax income share by 0.15 percentage points, or augment the percentage point change in income concentration by about 13%. We find similar reductions in inequality from assigning taxpayers with null-preparer prepares with tax optimization fixed effects two standard deviations above the mean (approximately the top 2% of preparers).

However, assigning top earners these same high-quality preparers would *increase* the top 1% post-tax income share by 0.4 percentage points, undermining the progressivity of the federal income tax by nearly 40%. This result highlights that although top earners disproportionately employ *better* tax preparers, there is substantial noise in this relationship: even among top 1% of earners, taxpayers do not strictly hire the best tax preparers, as illustrated by this hypothetical tax savings. This finding is further corroborated by the result that assigning null-preparer users the mean paid preparer of the top 1% of pre-tax earners yields almost no change to the top 1% income share.

5 Tax preparation fees

Up to this point, we have not discussed in detail the fees that preparers charge their clients. Consequently, the tax optimization that we estimate in Section 4 does not include the fees and represents instead what we refer to as gross tax savings. We make this choice for two reasons. First, we observe information on fees in our data only between 2012 and 2017. After the passage of the Tax Cuts and Jobs Acts (TCJA), itemizing taxpayers could no longer deduct tax preparation fees. Prior to 2012, our data on the tax preparation fee amount appears to be incomplete. Second, prior to TCJA this information is only available for those individuals who chose to itemize their deductions. Of those itemizers, only those whose "Job Expenses and Certain Miscellaneous Deductions" which exceed 2% of their AGI receive a benefit from reporting their tax preparation fees. Because of these data characteristics, we believe the observation of fees is highly skewed towards high income taxpayers and those with relatively high fees. For paid preparer users, this field corresponds with the fees paid to the tax preparers; for non-users, these fees likely correspond with fees paid for online tax preparation software assistance, such as TurboTax or H&R Block. Taxpayers who do not itemize their deductions take the standard deduction, a pre-determined amount based on the year and the taxpayer's filing-status. For those taxpayers, we do not observe information on fees. As such, while it is important to consider tax optimization inclusive of fees, the highly selected nature of the sample means that we instead consider fees separately.

5.1 Data and key concepts

We start by showing the selection associated with our sample reporting fees. In Figure A.4, we distinguish between preparer users and non-users. Panel (a) shows that we observe fees for around 20% of the population of taxpayers, while 40% of taxpayers in the population have a paid tax preparer but no observed fee information and 40% have no paid tax preparer and no fees.[29] We attribute our censoring to the fact that most taxpayers do not itemize their deductions. In fact, aggregate statistics from IRS SOI report that in 2015 29.6% returns featured itemized deductions, and higher-earners are disproportionately represented among itemizers, with 7.2% of returns of between $10,000 and $20,000 itemizing their deductions, but 96.8% of returns greater than 10 million USD itemizing. In our context, that means that the fees we observe are systematically selected on earnings (see Table 1 and

[29] Of the 20% for whom we do observe fees, 1/3 are taxpayers who do not report to have a tax preparer. While online software is nominally free, such services often charge a fee to users with "more complicated" return types, such as filing taxes in multiple US states, reporting capital income, and itemizing deductions.

Figure A.4 Panels (c) and (d)).

Since we are unaware of any other setting where a researcher can observe return-level fees, we proceed with empirical analysis in which we keep only individuals with non-missing fees. For this subgroup, we produce results that illustrate the relationship between fees and preparer quality and tax savings.[30]

Concept: net tax savings. As our measure of gross tax savings in Equation (5) does not take preparer fees into account, we define a new object of interest, *net* tax savings, where we net out the change in fees ϕ paid:

$$\begin{aligned}\text{Net preparer tax savings}_{it,J(it)} &= (B_{it} \cdot z_{it,j(it)} - \phi_{it,j(it)}) - (B_{it} \cdot z_{it,j=0} - \bar{\phi}_{it,j=0}) \\ &= \text{Gross tax savings}_{it,j(it)} - (\phi_{it,j(it)} - \bar{\phi}_{it,j=0}).\end{aligned} \tag{6}$$

The data limitations described above do not allow us to estimate the fees for taxpayers with the null-preparer using our mobility-identification strategy. Instead, we treat the null-preparer fee as constant, equal to the mean fee we observe among null-preparer users. While this decision means that we simply subtract a constant of $120 USD (2019), this feature is important for drawing inference on net tax positions with respect to preparer usage.[31] With this, in our analysis we include all taxpayers for whom we observe fees, both if they have a tax preparer or not. Further, taxpayers who use the null-preparer, desirably, do not mechanically have negative net preparer tax savings (as would be the case if one set their counterfactual fee to zero). Rather, while their gross preparer tax benefit of switching to the null-preparer is zero (because they already use the null-preparer), they may be assigned a net benefit if their empirically-observed fee is lower than the mean fee. This specification corresponds with the counterfactual exercise of moving empirically-observed null-preparer users to the "typical" null-preparer, including both measures of gross preparer tax savings and preparer costs.

[30] Table A.3 presents summary statistics for the fees data, pooling both non-users and users.

[31] Figure A.5 Panel (d) shows that average fees are stable over time and preparer-users on average pay USD 400.

5.2 Results

With these caveats in mind, we ask three questions: 1) do greater fees generate greater gross tax benefits, 2) what share of revenues preparers absorb from their gross tax benefits, and 3) how our measures of preparer optimization (preparer quality) correlate with fees. We examine these in turn in Figure 6.

In panel (a), we plot the distribution of fees conditional on income rank. Average fee is constant, around $200, until the 80th percentile of the income distribution. At this point, fees increase sharply up to around $1000 at the top 1% of the income distribution.[32]. Panel (b) plots the relationship between gross tax savings and fees in levels. The estimated slope of this relationship is about 2.3: for each Dollar increase in preparer fee, gross tax savings increase by more than 2.3 Dollars. This finding suggests that tax preparers take half of the tax savings they generate as own revenue. In columns (1)-(4) of Table 5, we quantify this relationship formally in logarithms. Column (1) shows that unconditionally, a 1% increase in gross tax savings generates a .235% increase in fees. Column (2) includes taxpayer fixed effects, which do not affect the estimate. Column (3) includes taxpayer-time varying covariates that reflect the complexity of the return (e.g. declaring business income, having other non-wage income, etc.), which absorb around 60% of the magnitude of the initial estimates. This result suggests that tax preparers who charge higher fees can generate higher gross savings partly because of more complex returns. Lastly, column (4) includes tax preparer fixed effects. The magnitude of the estimated coefficient drops to 0.006 and, although still statistically significant, suggests that most variation in how much gross tax savings a preparer can generate for a given fee is explained by variation *across* preparers rather than *within* preparer.

In Panel (c) of Figure 6 we calculate the fee as a share of gross tax savings, and plot this ratio against gross income rank. We show that the returns to using paid preparers increase for richer taxpayers.[33] While taxpayers at the bottom end of income distribution spend over

[32]Figure A.6 Panel (a) shows that this figure rises to $2000 among paid preparer users.

[33]Figure A.5 Panel (b) and Figure A.6 Panel (c) corroborate this result. The former plots fees as a share of broad income, showing that tax preparation fees evolve regressively in income rank, with paid preparer users in the bottom 20th percentile of broad earnings paying slightly above 4% of the pre-tax earnings on fees, decreasing to less than 0.5% of pre-tax earnings among the top fifth of earners. However, this relationship may be driven by some of the observed selection patterns on income. Figure A.6 Panel (a) plots a binned

60% of their gross tax saving on tax preparation fees, taxpayers at the top end of income distribution spend less than 20% of their gross tax saving on fees.

In Panel (d) we plot the relationship between preparer fixed effects and log fees paid. The figure shows that preparer quality, as measured by their ETR difference fixed effect, increases monotonically in fee. These results suggest that better preparers tend to charge their clients more. In column (5) of Table 5 we quantify this relationship, finding that a 1 log point increase in fees (an approximate tripling of fees) decreases the taxpayers effective tax rate by an additional one-fifth of a percentage point. In column (6), we include taxpayer fixed effects and find that although the slope is still significantly positive, the effect decreases by around 85%.

Finally, we consider how the results from Figure 4 change when we account for fees. First in Panel (e), we plot both average net and average gross tax savings as a share of broad income conditional on income rank. We find a substantial divergence between gross and net tax savings as a share of income, especially for lower earners, i.e., net preparer tax savings enter regressively along income rank, constituting a larger share of broad income, as income increases. However, this observation may be attributable to negative selection into itemization at the bottom of the income distribution: the taxpayers for whom we observe fees at the bottom of the income distribution may be using worse tax preparers (e.g. by merit of the observation that they are itemizing their deductions). Panel (f) breaks out the top 1% of earners and finds this relationship continues within the very top earners. These findings suggest that the savings that tax preparers offer to their clients, net of fees, reduce the progressivity of the income tax in the US.

6 Tax *optimization* or tax *evasion*?

In the final section, we investigate to what extent our measures of preparer tax optimization reflect tax evasion. We augment our data with information on audits and audit adjustments. Audits are very costly from the tax administration's perspective and taxpayers

scatter plot of log *net* tax savings against income rank, showing a positive relationship.

often use sophisticated methods of concealing under-reported income to elude audits (Guyton et al., 2023). In practice, fewer than 1% of taxpayers see an audit every year, and such audits (called "operational audits") are targeted toward taxpayers who 1. are suspected to not be in compliance, and 2. would yield a substantial increase in revenues to the tax authorities upon adjustment. Indeed, 85% of operationally audited tax returns see adjustment (IRS Statistics of Income, 2024). Furthermore, a tax authority would be less likely to audit a taxpayer highly suspected under-reporting a very small amount, as a successful audit would be costly but generate little gross return. Such gross returns are likely highly correlated with taxpayer true income.[34] Therefore, observed evasion is highly selected. While the National Research Program engages in randomized audits for research purposes, this would yield too small a sample to estimate preparer-level evasion fixed effects[35]

In practice, we use a binary variable for operational audit selection as an outcome variable and estimate Equation (2).[36] By definition, this variable exists for the universe of taxpayers, but its use requires careful interpretation. Given the selected nature of the sample, we control for income in our estimation procedure. We therefore interpret estimates from these audit selection regressions to represent suspected evasion and tax revenue returns to audit, conditional on income.

6.1 Results

Table 6 reports the results of this estimation procedure. The standard deviation of the preparer audit selection fixed effect is 0.5 percentage points. Figure 7 Panels (a)-(d) show

[34]Indeed, Figure A.7 Panels (a) and (b) plot the relationship between audit probability and broad income rank; audit probability is near zero until the fortieth percentile of broad earnings. Between p40 and p95, audit probability increases approximately linearly in income rank to around 0.5 percent. Audit probability sharply increases at this point, reaching 2 percent within the top 1 percentile of broad income.

[35]We are currently combining randomized audits and operational audits to further understand the role of evasion in our setting.

[36]In the ideal scenario, our data would feature perfect observation of tax non-compliance on the universe of taxpayers. In this case, we would estimate Equation (2) using either a binary variable for declaration-level tax noncompliance or another parameterization for the amount of noncompliance. This procedure would generate "tax evasion" fixed effects on the preparer level, which we would correlate with our tax preparer tax optimization fixed effects. The tax preparer evasion fixed effects would not imply a causal relationship between certain preparers and tax evasion (DeBacker et al., 2024), as tax preparers are not randomly assigned. The values of such fixed effects would indicate either sorting of evading taxpayers sort into specific preparers or indeed that preparers cause evasion.

conditional distributions of the audit probability fixed effect estimates. Panel (a) plots the relationship between tax preparer audit fixed effect and taxpayer income rank. It shows that richer taxpayers hire preparers that reduce their chance of audit by up to nearly 0.3 percentage points among the top percentile. The figure also shows a relative reduction in audit probability among the preparers hired by the poorest taxpayers, who likely hire tax preparers in assistance with credit filing. Table 7 columns (1) - (4) quantify this relationship but suggest that taxpayer income rank has very little explanatory power in describing the allocation of tax preparer audit probability fixed effects across taxpayers. Panel (b) shows the relationship between taxpayer audit probability fixed effect and income rank; the figure depicts a U-shape with a right skew and a flat, near-zero interior. Note that the poorest taxpayers see mildly outsized audit probability, as such taxpayers are more likely to claim refundable credits which may trigger an audit. Panel (c) depicts the negative sorting between tax preparers with lower audit probability among taxpayers with higher audit probability: this assortative matching is intuitive, as taxpayers with higher audit probability likely hire tax preparers that will decrease their probability of audit, as taxpayers likely view audits as an undesirable outcome.

Figure 7 Panel (d) shows the preparer-level correlation between preparer audit probability fixed effect and preparer tax optimization fixed effect. The figure depicts a monotonically increasing relationship between the two variables with a slope of about 0.2: a one percentage point increase in preparer ETR difference fixed effect increases the preparer audit probability fixed effect by 0.2.[37] Table 7 columns (5) - (6) show that tax preparer optimization explains around 2 percent of the cross-preparer variation in audit probability fixed effect. This finding suggests that greater tax optimization is associated with greater frequency of non-random observation of tax noncompliance. While this is striking, it may be mechanically the case that greater tax optimization is generated by taking more aggressive tax positions whose compliance may not hold up to scrutiny in the case of an audit.

[37]Figure A.8 displays versions of this panel using an alternate measure of tax optimization as well as using analytic weighting on the square-root of number of observations used to construct each preparer fixed effect; the results are similar. However, with the large mass of individuals using the null-preparer, there is a discontinuity at the origin in these figures corresponding with the null-preparer.

7 Conclusion

The role of income taxation is both to raise revenue for public expenditures and to regulate inequality by converting pre-tax income into post-tax income. However, paid tax professionals have largely been overlooked in the central role they play in mediating this process. While over half of Americans use paid tax preparers, tax preparers are also important in countries with automatic income tax filing, where their use is likely even more concentrated among top earners.

We document the revenue and distributional impacts of tax preparers using data from the universe of US individual income taxpayers. We develop a novel approach for measuring declaration-level tax optimization and apply this measure in a mobility-based design to estimate the causal impacts of tax preparers on tax-optimization.

We find that tax preparers reduce tax revenue collections and the progressivity of the US federal income tax. Moving to a preparer of one standard deviation better quality reduces one's effective income tax rate by 0.5 percentage points. However, we document substantial heterogeneity in tax preparer quality by income. The top percentile of earners hire preparers that reduce their ETRs by around 0.7 percentage point on average, whereas and bottom-earners benefit from only a 0.2 percentage point reduction in ETR from their preparers. We conduct simple benchmarking exercises and find that assigning the top 1% of pre-tax earners to the outside option of not hiring a paid preparer would augment the inequality-reducing properties of the personal income tax system by 13%.

We further explore paid preparer usage by studying the fees they charge their clients. We find that on average, preparers extract around half of their client's gross tax savings as revenues. Fees charged correlate positively with tax preparer quality, explaining 11% in the variation in the preparer fixed effect; however, we find that nearly all of the variation in fees occurs *across* preparers, rather than within preparers charging more or less for different tax services. We also find that taking fees into account, paid preparer usage further undermines the progressivity of the US income tax schedule, a net tax savings as a share of income are on average negative for low earners and are increasing in income rank.

Lastly, we investigate the relationship between tax preparer optimization fixed effects and audit selection. We use audit selection as a measure indicating suspected underreporting and expected returns to audit. 85% of operational audits indeed see some adjustment indicating tax "evasion", so we interpret this variable to indicate high likelihood of tax non-compliance (conditional on income). We find a significant positive relationship between preparer-level audit probabilities and tax optimization, indicating that the preparers that are better at reducing their clients' tax burdens are also more frequently suspected of generating tax noncompliance. While this result does not suggest a causal relationship between specific tax preparers and tax evasion, the finding shows that tax preparers who reduce their client's tax burdens more are suspected more frequently of tax noncompliance and are indeed selectively observed more frequently as tax non-compliant.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Alstadsæter, A., Johannesen, N., and Zucman, G. (2018). Who owns the wealth in tax havens? macro evidence and implications for global inequality. *Journal of Public Economics*, 162:89–100. In Honor of Sir Tony Atkinson (1944-2017).

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

Auten, G. and Splinter, D. (Forthcoming). Income inequality in the united states: Using tax data to measure long-term trends. *Journal of Political Economy*.

Ayres, F. L., Jackson, B. R., and Hite, P. S. (1989). The economic benefits of regulation: Evidence from professional tax preparers. *Accounting Review*, pages 300–312.

Battaglini, M., Guiso, L., Lacava, C., and Patacchini, E. (2022). Tax professionals and tax evasion. *Working Paper*.

Belnap, A., Hoopes, J. L., and Wilde, J. H. (2024). Who really matters in corporate tax? *Journal of Accounting and Economics*, 77(1):101609.

Boccanfuso, J. and Ferey, A. (2023). Inattention and the taxation bias. *Journal of the European Economic Association*, 22(3):1452–1494.

Bonhomme, S., Holzheu, K., Lamadon, T., Manresa, E., Mogstad, M., and Setzler, B. (2023). How much should we trust estimates of firm effects and worker sorting? *Journal of Labor Economics*, 41(2):291–322.

Bonhomme, S., Lamadon, T., and Manresa, E. (2019). A distributional framework for matched employer employee data. *Econometrica*, 87(3):699–739.

Bonhomme, S., Lamadon, T., and Manresa, E. (2022). Discretizing unobserved heterogeneity. *Econometrica*, 90(2):625–643.

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

Boning, W. C., Hendren, N., Sprung-Keyser, B., and Stuart, E. (2025). A welfare analysis of tax audits across the income distribution. *The Quarterly Journal of Economics*, 140(1):63–112.

Breunig, R., Deutscher, N., and Hamilton, S. (2024). Rounded up: Using round numbers to identify tax evasion. *Journal of Public Economics*, 238:105195.

Brounstein, J. (2023). Can countries unilaterally mitigate tax haven usage? evidence from an ecuadorian transaction tax. *Working paper*.

Bureau, U. C. (2022). Professional, scientific, and technical services: Summary statistics for the u.s., states, and selected geographies: 2022ëconomic census, ecn core statistics summary statistics for the u.s., states, and selected geographies: 2022, table ec2254basic.

Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1):S13–S70.

Card, D., Cardoso, A. R., and Kline, P. (2015). Bargaining, Sorting, and the Gender Wage Gap: Quantifying the Impact of Firms on the Relative Pay of Women. *The Quarterly Journal of Economics*, 131(2):633–686.

Card, D., Cardoso, A. R., and Kline, P. (2016). Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women. *The Quarterly journal of economics*, 131(2):633–686.

Card, D., Heining, J., and Kline, P. (2013a). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics*, 128(3):967–1015.

Card, D., Heining, J., and Kline, P. (2013b). Workplace Heterogeneity and the Rise of West German Wage Inequality*. *The Quarterly Journal of Economics*, 128(3):967–1015.

Chetty, R., Friedman, J. N., and Saez, E. (2013). Using differences in knowledge across neighborhoods to uncover the impacts of the eitc on earnings. *American Economic Review*, 103(7):2683–2721.

Christian, C. W., Gupta, S., and Lin, S. (1993). Determinants of tax preparer usage: Evidence from panel data. *National Tax Journal*, 46(4):487–503.

DeBacker, J., Heim, B. T., Tran, A., and Yuskavage, A. (2024). Seeking professional help: How paid preparers decrease tax compliance. *National Tax Journal*, 77(2):229–261.

Feenberg, D. R. and Coutts, E. (1993). An introduction to the taxsim model. *Journal of Policy Analysis and Management*, 12(1):189–194.

Finkelstein, A., Gentzkow, M., and Williams, H. (2021). Place-based drivers of mortality: Evidence from migration. *American Economic Review*, 111(8):2697–2735.

Goldin, J., Homonoff, T., Javaid, R., and Schafer, B. (2022). Tax filing and take-up: Experimental evidence on tax preparation outreach and benefit claiming. *Journal of Public Economics*, 206:104550.

Guyton, J., Langetieg, P., Reck, D., Risch, M., and Zucman, G. (2023). Tax evasion at the top of the income distribution: Theory and evidence. (28542).

IRS Statistics of Income (2024). Soi tax stats - examination coverage and recommended additional tax after examination, by type and size of return - irs data book table 17.

Kaçamak, Y. (2022). The role of information aggregators in tax compliance. Technical report.

Kleven, H., Kreiner, C., Larsen, K., and Søgaard, J. (2024). Micro vs macro labor supply elasticities: The role of dynamic returns to effort. *Working Paper*.

Kline, P., Saggio, R., and Sølvsten, M. (2020). Leave-out estimation of variance components. *Econometrica*, 88(5):1859–1898.

Kopczuk, W. and Pop-Eleches, C. (2007). Electronic filing, tax preparers and participation in the earned income tax credit. *Journal of public Economics*, 91(7-8):1351–1367.

Londoño-Vélez, J. and Ávila Mahecha, J. (2021). Enforcing wealth taxes in the developing world: Quasi-experimental evidence from colombia. *American Economic Review: Insights*, 3(2):131–48.

Mincer, J. (1958). Investment in human capital and personal income distribution. *Journal of Political Economy*, 66(4):281–302.

Piketty, T. and Saez, E. (2003). Income inequality in the united states, 1913–1998. *The Quarterly Journal of Economics*, 118(1):1–41.

Piketty, T., Saez, E., and Zucman, G. (2017). Distributional National Accounts: Methods and Estimates for the United States*. *The Quarterly Journal of Economics*, 133(2):553–609.

ProPublica (2019). Inside turbotax's 20-year fight to stop americans from filing their taxes for free. *ProPublica*.

Reidenbach, M., Sorensen, T. L., and Treu, J. S. (2021). Regulation and tax preparer qualifications. *The Journal of the American Taxation Association*, 43(1):125–147.

Saez, E. and Zucman, G. (2016). Wealth inequality in the united states since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, 131(2):519–578.

Smith, M., Zidar, O., and Zwick, E. (2022). Top wealth in america: New estimates under heterogeneous returns. *The Quarterly Journal of Economics*, 138(1):515–573.

Tortarolo, D., Cruces, G., and Castillo, V. (2020). It takes two to tango: Labor responses to an income tax holiday in argentina. *Working Paper*.

Treasury Inspector General for Tax Administration (2020). Complexity and insufficient oversight of the free file program result in low taxpayer participation.

Zwick, E. (2021). The costs of corporate tax complexity. *American Economic Journal: Economic Policy*, 13(2):467–500.

Main figures and tables

Table 1: Descriptive characteristics of taxpayers and tax preparers
Panel (a): Taxpayer characteristics

	Non-users		Preparer users		Difference			
	Mean	SD	Mean	SD	Unsaturated	Controlling for log broad income	Taxpayer fixed effects	N
Age	42.8	16.6	50.3	17.6	7.56	6.04	.085	956328832
Female	.365	.481	.282	.45	-.083	-.044	-	958975488
Broad income	91318	20272839	157624	10714834	66306	-28299	11982	959623488
Log broad income	10.9	.864	11.2	1	.322	0	.062	959623488
AGI	75016	418802	118551	1043516	43535	-5046	9956	959623488
Log AGI	10.9	.849	11.1	.956	.202	-.083	.026	959623488
Taxable income	59781	1725357	91918	1254152	32137	-10303	5336	959623488
Log taxable income	10.3	1.25	10.6	1.34	.218	-.177	.008	914841344
Taxes paid	8930	134422	19524	290711	10595	-783	2475	959623488
Log taxes paid	8.19	1.49	8.49	1.62	.292	-.171	.028	910039680
Tax savings	10047	7743079	20533	4015601	10486	-10105	721	959623488
Log tax savings	8.2	.965	8.65	1.23	.451	.122	.098	959622784
ETR	.077	.05	.078	.057	.001	-.01	-.001	959623488
ETR difference	.074	.035	.087	.048	.014	.011	.004	959623488
Prop. tax optimization	.516	.247	.549	.259	.033	.063	.014	959623488
Unique taxpayers	100151511		111417835					
Total unique taxpayers	175635241							

Panel (b): Preparer characteristics

	Mean	SD	p1	p10	p25	p50	p75	p90	p99	N
Unique clients	156	333	1	1	3	30	158	460	1112	1267866
Unique clients per year	63.3	125	1	1	2.51	18.2	70.7	174	424	1267866
Probability client stays next year	.373	.368	0	0	0	.275	.753	.902	.961	1166559
Year of birth	1967	16.4	1932	1946	1954	1967	1980	1988	1993	1038399
Female	.353	.469	0	0	0	0	1	1	1	1038281
Years active in panel	4.58	3.65	1	1	1	3	9	10	10	1267866
Is attorney	.017	.129	0	0	0	0	0	0	.6	1267866
Is CPA	.174	.378	0	0	0	0	0	1	1	1267866
Fee	481	1606	62	139	200	277	391	730	2982	656179
Preparer AGI	96634	218045	-132	12274	24245	58013	119140	208865	484832	1038467
Mean client log broad income	11	.737	9.57	10.3	10.5	10.9	11.4	11.9	12.8	1267866
Mean client ETR	.071	.034	0	.035	.052	.067	.085	.11	.158	1267866
Mean client ETR difference	.087	.028	.032	.062	.073	.083	.097	.114	.165	1267866
Mean client prop. tax optimization	.573	.147	.174	.416	.498	.562	.638	.754	.901	1267866
Unique preparers	1267866									

Note: This table shows descriptive statistics of taxpayers and tax preparers. Panel (a) displays means and standard deviations for various dependent variables on the taxpayer-year level, stratifying by paid preparer status. The three columns under "difference" display univariate regression coefficients for the difference between preparer users and non-users, with the specification labeled accordingly. Standard errors are clustered on the taxpayer level are omitted for brevity, as all coefficients correspond with a t-value of at least 10, except for the coefficient under "Taxpayer fixed effects" corresponding with "tax savings", which has a t-statistic of 1.67. The N columns corresponds with the number of observations in the unsaturated regression. Panel (b) displays distributional statistics for tax preparers. For confidentiality purposes, percentile values correspond with fuzzy percentiles computed as the mean of the five surrounding .1 percentile values (e.g. p50 is computed as the mean of p49.8, p49.9, p50, p50.1, and p50.2).

Figure 1: Tax preparer usage and users throughout the income distribution

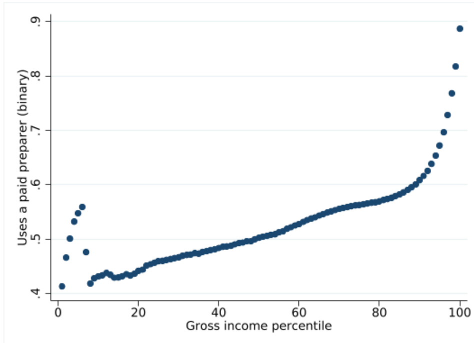

Panel (a): Whether a taxpayer uses a paid preparer

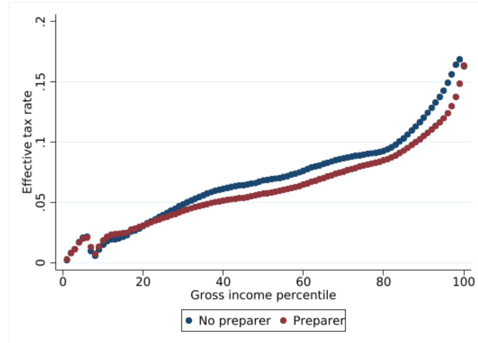

Panel (b): Effective tax rate

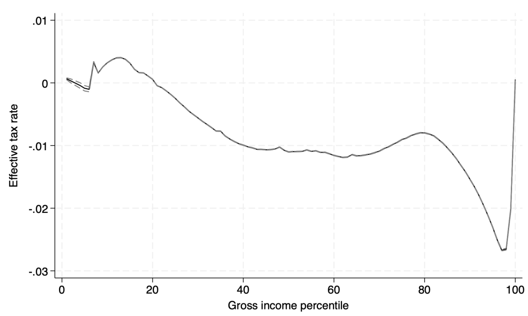

Panel (c): Difference in effective tax rate
Preparer users v. non-users

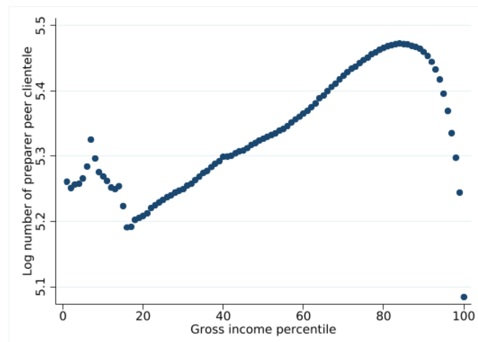

Panel (d): Log number of preparer peer-clientele

Note: These figures display binned scatter plots of various characteristics of paid preparer usage along gross income percentile. Gross income percentiles are constructed as within-year percentile ranks of tax-filing units, where gross income represents income across different reported sources within the income tax declaration prior to the application of credits and deductions. Panel (c) plots the coefficients from a regression of effective tax rate on the interaction of binary paid preparer usage and income percentile, with standard errors clustered on the individual level. The dependent variable in Panel (d) corresponds with the log number of individuals who share an identical tax preparer in the same year.

Figure 2: The characteristics of tax optimization measure "ETR difference"

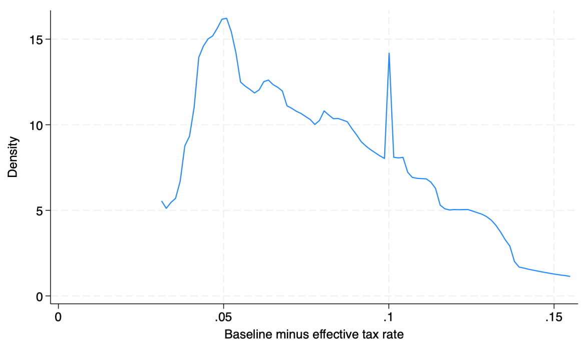

Panel (a): distribution of ETR difference

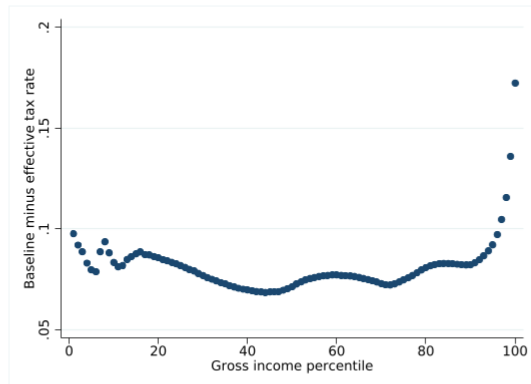

Panel (b) ETR difference along the broad income distribution

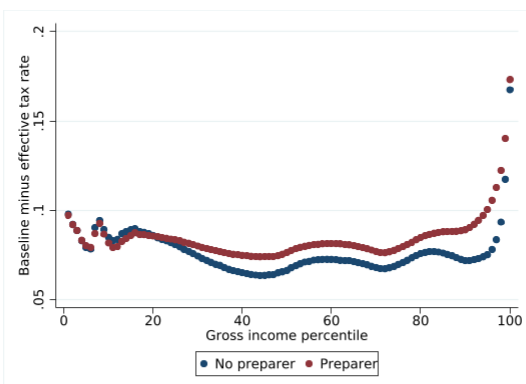

Panel (c): ETR difference along the broad income distribution: Preparer users v. non-users

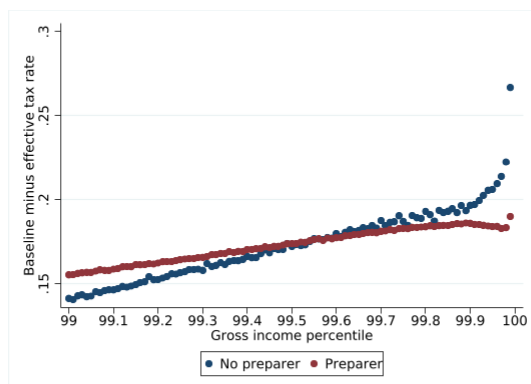

Panel (d) ETR difference within the top 1% of broad income: Preparer users v. non-users

Note: Panel a shows a histogram of the distribution of ETR difference for all taxpayer-year observations in our data. Panels c - d show binned scatter plots of ETR difference measure across the percentiles of income distribution. ETR difference is defined as the percentage point difference in baseline ETR as imputed statutorily from our broad income concept and final, realized ETR. Gross income percentiles are constructed as within-year percentile ranks of tax-filing units, where gross income represents income across different reported sources within the income tax declaration prior to the application of credits and deductions.

Figure 3: Evolution of taxpayer ETR difference around preparer moves:
By incumbent and destination preparer leave-out mean ETR difference

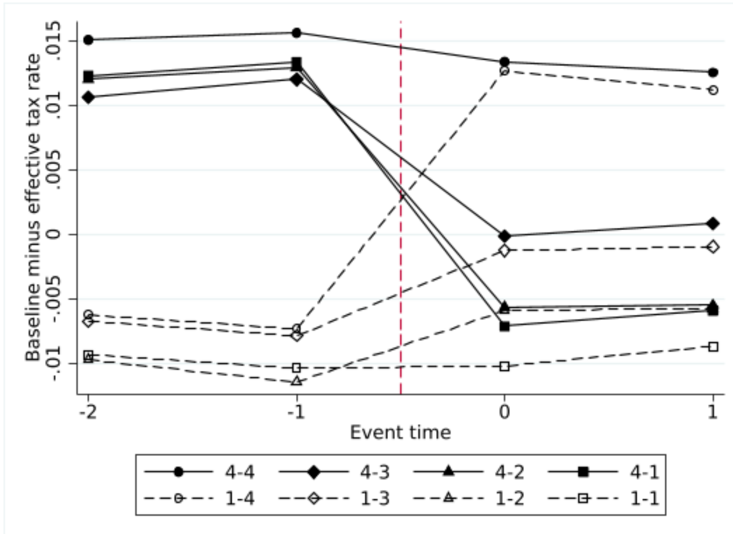

Panel (a): ETR Difference

Note: This figure shows mean levels of ETR difference for taxpayers that switch tax preparers. We consider the set of taxpayer-preparer switches in which the incumbent and incoming match lasted at least two years. Each preparer is classified into quartiles based on the mean value of ETR difference of preparer peer-clientele in the same year. Within each move-type, we plot the average ETR difference by event-time. The dependent variable is constructed as the percentage point difference between ETR based on statutory tax imputed from broad income and final realized ETR, which is residualized on year and state fixed effects and log income.

Table 2: Tax preparer and payer heterogeneity in tax optimization

			By Income		
	All	No null preparers	Low	Medium	High
	(1)	(2)	(3)	(4)	(5)
(i) ETR diff					
STD preparers (ψ)	0.005	0.006	0.003	0.005	0.006
STD taxpayers (α)	0.027	0.032	0.017	0.021	0.038
$corr(\psi, \alpha)$	0.087	0.007	0.054	0.066	0.039
# of preparers	1,267,867	1,267,866	1,068,091	1,049,060	947,050
# of taxpayers	171,237,396	111,417,835	78,770,676	51,150,655	41,316,065
(ii) Deductions					
STD preparers (ψ)	0.027	0.029	0.021	0.027	0.025
STD taxpayers (α)	0.157	0.175	0.136	0.144	0.169
$corr(\psi, \alpha)$	0.169	0.063	0.046	0.128	0.148
# of preparers	1,267,867	1,267,866	1,068,091	1,049,060	947,050
# of taxpayers	171,237,396	111,417,835	78,770,676	51,150,655	41,316,065
(iii) Credits					
STD preparers (ψ)	0.012	0.016	0.013	0.012	0.015
STD taxpayers (α)	0.074	0.077	0.058	0.088	0.062
$corr(\psi, \alpha)$	-0.020	0.008	0.008	-0.049	-0.005
# of preparers	1,267,867	1,267,866	1,068,091	1,049,060	947,050
# of taxpayers	171,237,396	111,417,835	78,770,676	51,150,655	41,316,065

Note: This table displays standard deviations and correlations between different fixed effect objects from our estimation procedure. Panel (i) contains estimates pertaining to our main tax optimization variable, ETR difference. Panels (ii) and (iii) present estimates that distinguish between optimization attributable to deductions and credits respectively. Column (1) displays estimates from our main specification. Column (2) contains estimates from an alternate specification that omits taxpayers without a paid tax preparer. Columns (3) - (5) stratify our estimation procedure by within-year broad income tercile. Each specification indicates the final number of unique tax preparers and tax payers that constitute the leave-one-out largest connected set.

Figure 4: Tax preparer fixed effects across the income distribution.

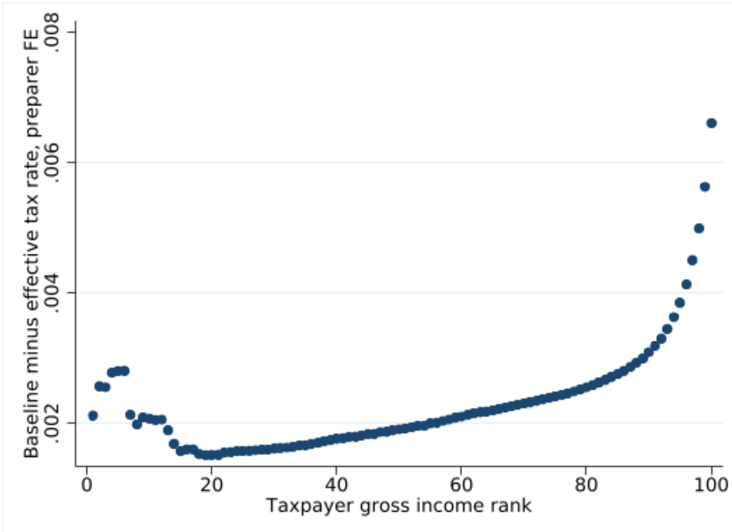

Note: This figure displays a binned scatter plot of tax preparer ETR difference fixed effect by broad income rank. Broad income percentiles are constructed within-year on the universe of taxpayers prior to introducing sample restrictions. ETR difference corresponds with the levels difference in effective tax rate imputed statutorily from broad income and true effective tax rate on final taxes paid.

Table 3: Correlations between tax preparer ETR difference fixed effects and tax preparer characteristics.

	Coefficient (pp)	Standard error (pp)	T-statistic	P-value	N	Adjusted R2
Unique clients	-.0000899	2.81e-06	-31.97927	0	985521	.000592
Unique clients per year	-.0000656	7.49e-06	-8.758345	1.98e-18	985521	.0000442
Probability client stays next year	-.289	.00377	-76.75098	0	956772	.0064
Year of birth	-.00691	.000203	-33.97256	0	912853	.00781
Female	-.0389	.00297	-13.0853	0	912743	.000199
College	-.0866	.00435	-19.89077	0	912911	.000512
Years active in panel	-.00481	.000369	-13.04312	7.01e-39	985521	.000164
Is attorney	.0641	.0141	4.552536	5.30e-06	985521	.0000359
Is CPA	.078	.00338	23.05216	0	985521	.000532
Log fee	.267	.00284	94.0118	0	641800	.0323
Preparer log AGI	-.00973	.00119	-8.161366	3.32e-16	907759	.0000882
Mean client log broad income	.135	.00319	42.32248	0	985521	.00397

Note: This table shows correlations between tax preparer fixed effects. Each row corresponds with a univariate cross-sectional regression of preparer-level ETR difference fixed effect on the respective row variable. The sample consists of the universe of paid preparers that have an estimated ETR difference fixed effect (excluding the null-preparer) where each observation is a unique paid preparer. The constant is omitted for legibility. The mean ETR difference fixed effect among preparers is 0.5 percentage points, and coefficients and standard errors are multiplied by 100 to represent percentage points. Standard errors are heteroskedasticty robust to heteroskedasticity

Figure 5: Counterfactual revenue benchmarks.

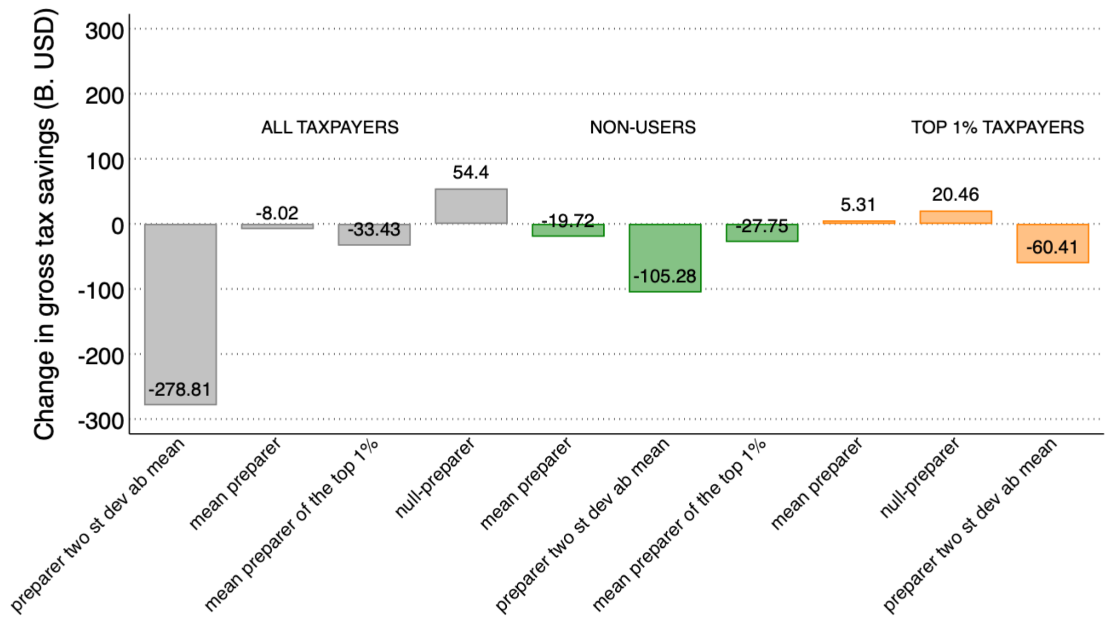

Note: This figure presents potential changes in aggregate gross tax savings from different sets of counterfactual exercises. In grey we mark counterfactuals where we assign all taxpayers different tax preparers. In green, we mark counterfactuals where we assign all non-users different tax preparers. In orange, we mark counterfactuals where we assign only top 1% of taxpayers different tax preparers. Numbers on top of each bar are the changes in aggregate gross tax savings from each experiment in billions of USD.

Table 4: Counterfactual inequality benchmarks for different allocations of tax preparers

Benchmarks:	
Top 1% pre-tax income share	.2427
Top 1% post-tax income share	.2312
Post- less pre-tax change in top 1% income share	**-.0115**
Changes relative to the post-tax top 1% income share:	
All taxpayers are assigned:	
The mean preparer	-.0007
The null-preparer	-.0008
The preparer two standard deviations above the mean	-.0004
The mean preparer of the top 1%	-.0007
The top 1% are assigned:	
The mean preparer	-.0004
The null-preparer	-.0015
The preparer two standard deviations above the mean	.0043
Non-users are assigned:	
The mean preparer	-.0002
The preparer two standard deviations above the mean	-.0011
The mean preparer of the top 1%	-.0003
Mean preparer ETR difference fixed effect	.0051

Note: This table displays select estimates for the top 1% post-tax income share under different scenarios of the allocation of tax preparers. Each line corresponds with a counterfactual scenario where we keep the broad income, taxpayer characteristics and year fixed, but replace the tax preparer fixed effect accordingly. The top 1% of earners by broad income is defined on the universe taxpayers prior to the introduction of sample restrictions.

Figure 6: Fees, tax savings, and preparer quality

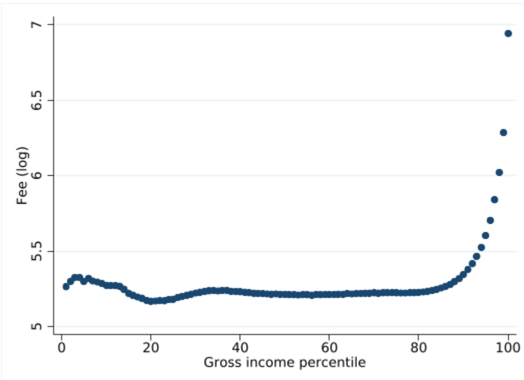

Panel (a): Log fees

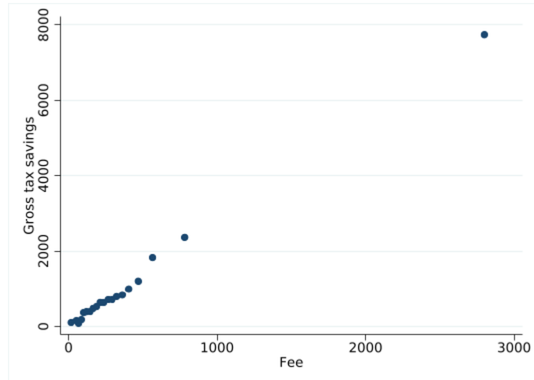

Panel (b): Gross tax savings per Dollar preparer fee

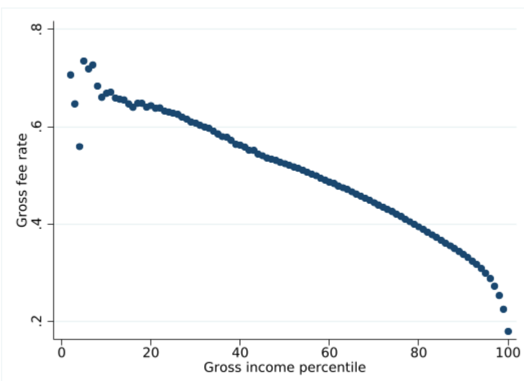

(c) Gross fee rate by income rank

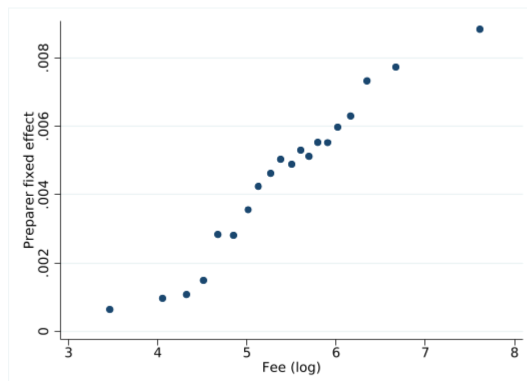

(d) Preparer quality by log fee

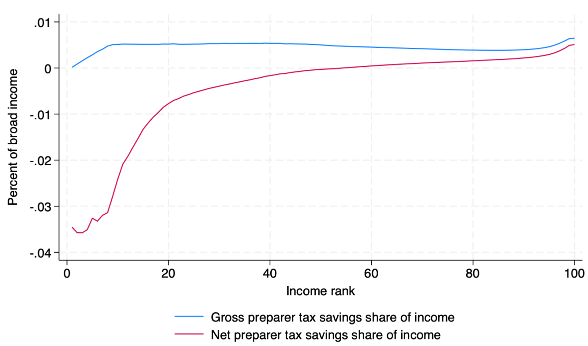

(e) Preparer net tax savings as a share of income
By income rank

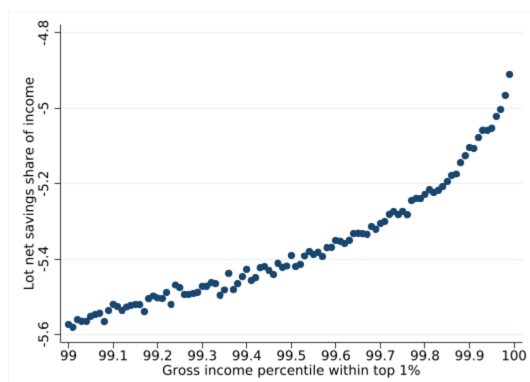

(f) Log net preparer tax savings as a share of income
By income rank
Within the top 1%

Note: These figures plot results pertaining to fees. Gross income percentiles are computed as within-year ranks of broad income among taxpayers. Fees are observed for the subsample of our data from 2012 to 2017 that itemized their deductions.

Table 5: Quantifying the relationship between fees, tax savings and preparer quality.

	Log fee				Preparer ETR difference fixed effect	
	(1)	(2)	(3)	(4)	(5)	(6)
Log gross preparer tax savings	.235	.234	.098	.006		
	(.002)	(.001)	(.002)	(0)		
Log fee					.00225	.00039
					(.00024)	(.00007)
Constant	4.1	4.1	0	5.6	-.007	.002
	(0)	(0)	(.2)	(0)	(.002)	(0)
N	64510867	64471727	64510867	56605445	101800000	89746749
No. clusters	458687	419547	458687	430166	604971	575968
Time-varying controls	No	No	Yes	No	No	No
Fixed effects	None	Taxpayer	None	Preparer	None	Taxpayer
Adj. R2	.1517	.5882	1.7024	.8846	.10844	.85464

Note: This table displays coefficients and statistics from a series of regressions of fees and preparer fixed effects. Fees are observed on a subset of our data from 2012-2017. Time-varying controls include individual-year values of log broad income and binary variables for whether the individual files for the EITC or has additional business income schedules C or E. Standard errors are clustered on the preparer-level.

Table 6: Tax preparer and preparer heterogeneity in audit probability

			By Income		
	All (1)	*No null-preparers* (2)	Low (3)	Medium (4)	High (5)
(i) Audited					
STD preparers (ψ)	0.005	0.007	0.003	0.006	0.007
# of preparers	1,267,868	1,267,867	1,068,871	1,048,529	945,942
# of taxpayers	175,635,241	119,085,635	79,082,823	51,007,354	41,147,219

Note: This table displays standard deviations and correlations between different fixed effect objects from our two-way fixed effects estimation procedure. The dependent variable in all specifications is a binary variable for whether an taxpayer was selected for an operational audit. Column (1) displays estimates from our main specification. Column (2) contains estimates from an alternate specification that omits taxpayers without a paid tax preparer. Columns (3) - (5) stratify our estimation procedure by within-year broad income tercile. Each specification indicates the final number of unique tax preparers and tax payers that constitute the leave-one-out largest connected set.

Table 7: Regressions of preparer audit fixed effects on taxpayer and preparer characteristics

	Taxpayer-year level				Preparer level			
	Income rank		Income rank within top 1%		ETR difference preparer fixed effect		Proportion tax optimization preparer fixed effect	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Coefficient	-.0000144	-2.54e-10	-9.44e-06	-1.26e-09	.239	.19	.0416	.0328
	(1.60e-08)	(3.46e-13)	(2.34e-07)	(1.12e-11)	(.00452)	(.0443)	(.00082)	(.0103)
Constant	.0000951	3.21e-09	-.00222	-6.60e-08	-.00137	-.00126	-.00143	-.00128
	(9.36e-07)	(1.67e-11)	(9.41e-06)	(3.95e-10)	(.0000269)	(.000676)	(.0000268)	(.000734)
Weight	Equal	$\sqrt{N_j}$	Equal	$\sqrt{N_j}$	Equal	$\sqrt{N_j}$	Equal	$\sqrt{N_j}$
N	959098310	959098310	13653055	13653055	985519	985519	985520	985520
R2	.00284	8.57e-08	.000659	5.85e-07	.0154	.0253	.0145	.0201

Note: This table displays univariate regression coefficients of preparer-level audit fixed effect estimates on different independent variables. Each column corresponds with a separate regressions, with the column title corresponding with the left-hand-side whose coefficient is estimated in the row indicated "Coefficient". Weights $\sqrt{N_j}$ correspond with analytic weights on the number of observations used to compute the fixed effect of preparer j. Columns (1)-(4) correspond with taxpayer-year level regressions with standard errors clustered on the taxpayer level. Columns (5)-(9) correspond with cross-sectional regressions on the unique tax preparer level with heteroskedasticity-robust standard errors estimated in parentheses.

Figure 7: Audit results: preparer audit fixed effects

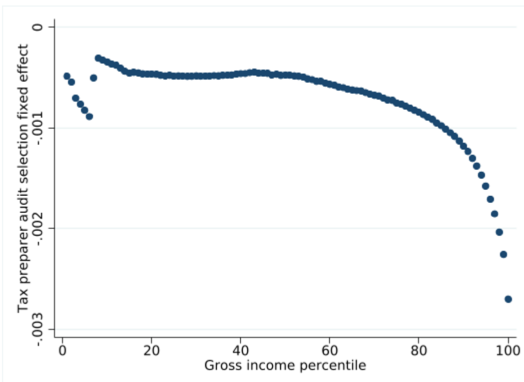

Panel (a): Preparer audit fixed effect conditional on taxpayer income rank

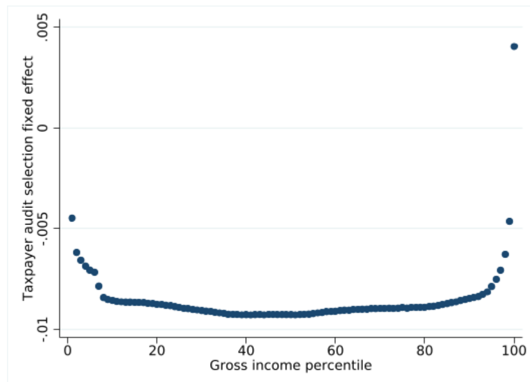

Panel (b): Taxpayer audit fixed effect rank by taxpayer income rank

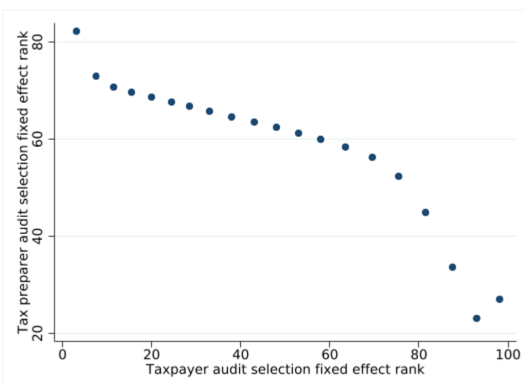

Panel (c): Preparer audit fixed effect rank versus taxpayer audit fixed effect rank

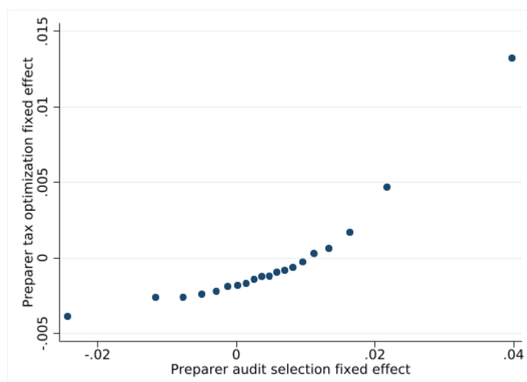

Panel (d): Preparer audit fixed effect versus ETR difference fixed effect

Note: These figures are binned scatter plots of preparer and taxpayer fixed effects and taxpayer income rank. Tax preparer and taxpayer fixed effects are computed using a two-way fixed effects regression with a binary variable for operational audit selection as the dependent variable. Taxpayer income rank is computed within-year. Fixed effect ranks are computed over the unique set of taxpayer or tax preparer fixed effects.

A Additional figures and tables

Table A.1: Taxpayer/preparer characteristics: main sample v. largest connected leave-one-out set

Panel (a): Taxpayer characteristics

	Baseline	Excluded	Connected set	Difference
Taxpayer-year observations	9.600e+08	1.550e+08	8.050e+08	
Unique taxpayers	1.710e+08	4.200e+07	1.290e+08	
Age	46.93	44.07	47.48	-3.418
	(17.575)	(16.57)	(17.708)	(.004)
Female	0.319	0.319	0.319	0
	(.466)	(.466)	(.466)	(0)
College	0.162	0.163	0.162	0.00100
	(.369)	(.37)	(.369)	(0)
Has paid preparer	0.548	0.620	0.534	0.0860
	(.498)	(.485)	(.499)	(0)
Gross income	127674	99805	133041	-33236
	(1.58e+07)	(1.02e+07)	(1.66e+07)	(1024.402)
Taxable income	77402	62680	80237	-17556
	(1485700)	(1411020)	(1499638)	(159.433)
Taxes paid	14739	10714	15514	-4800
	(233514.1)	(166581.7)	(244299.3)	(33.1)
ETR	0.0780	0.0740	0.0780	-0.00400
	(.054)	(.05)	(.054)	(0)
Baseline minus effective tax rate	0.0810	0.0780	0.0820	-0.00400
	(.043)	(.038)	(.044)	(0)
Final tax optimization	0.534	0.535	0.534	0
	(.254)	(.248)	(.255)	(0)

Panel (b): Preparer characteristics

	Baseline	Excluded	Connected set	Difference
Taxpayer-year observations	9.600e+08	514268	9.590e+08	
Unique taxpreparers	1.268e+06	280336	987531	
Preparer characteristics				
Unique clientele per year	77.91	1.220	99.68	-98.46
	(35783.9)	(.772)	(40546.09)	(40.801)
Proportion of frame active	0.458	0.139	0.549	-0.410
	(.365)	(.123)	(.36)	(0)
Clientele continuation probability	0.524	0.907	0.490	0.417
	(.286)	(.207)	(.266)	(.001)
Preparer AGI	96634	98404	96393	2011
	(218044.8)	(250739.8)	(213216.6)	(746.343)
Clientele characteristics				
Gross income	147866	190340	135809	54531
	(3250400)	(4021765)	(2995339)	(8204.936)
ETR	0.0690	0.0700	0.0690	0.00100
	(.054)	(.056)	(.053)	(0)
Final tax optimization	0.586	0.589	0.585	0.00400
	(.259)	(.264)	(.258)	(.001)

Note: This table shows descriptive statistics of taxpayers and tax preparers based on whether they are included in the leave-one-out largest connected set. The Connected Set column constitutes are main sample of analysis.

Figure A.1: Time series of preparer usage

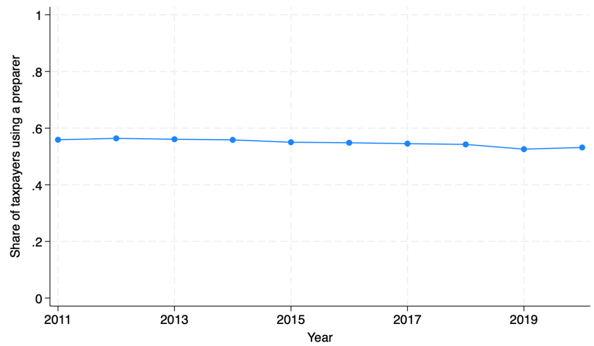

Panel (a): Share of taxpayers with a preparer

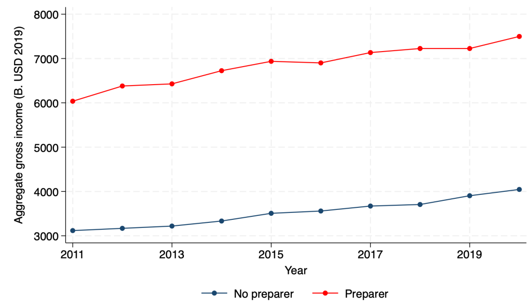

Panel (b): Aggregate fiscal income by preparer usage

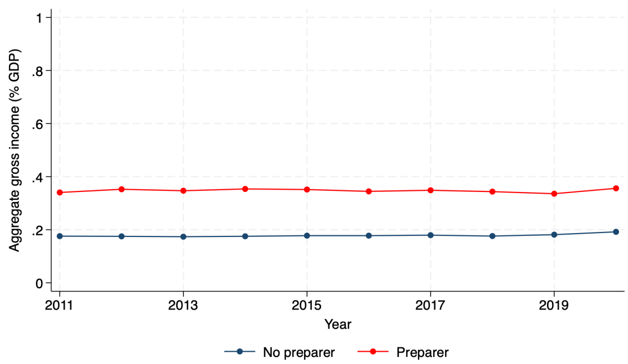

Panel (c): Fiscal income by preparer usage, share of GDP

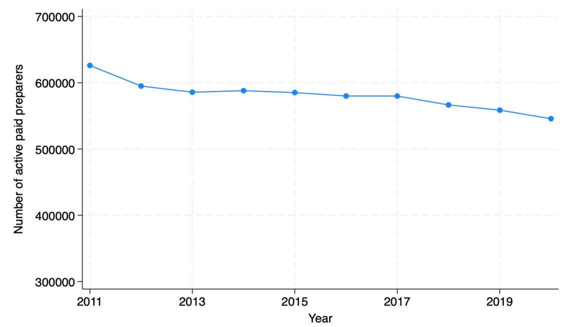

Panel (d): Time series of number of preparers

is constructed as the number of unique paid preparers observed every year.

Note: This figure shows time series of around preparer usage. Panel (c) plots the evolution in the share of GDP allocable to taxpayers with and without paid preparers; note that each year's shares do not add up to one because GDP exceeds fiscal income (the income reported to the tax authorities by individual taxpayers). Number of active paid preparers in Panel (d)

Figure A.2: Probability of using a paid tax preparer by income rank
Wager versus non-wage earners

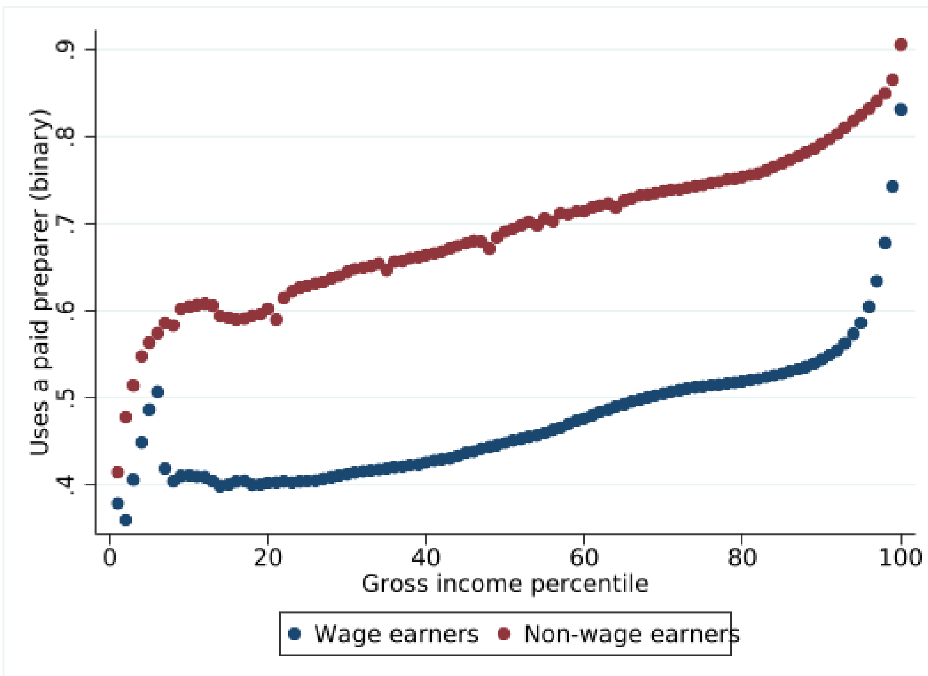

Note: This figure is a binned scatter plot of whether a taxpayer uses a paid tax preparer by gross income rank. The estimation is stratified based on whether one is a wage earner or non-wage earner, where wage earners report at least half of their broad income as ordinary wage income from a form W-2.

Figure A.3: Additional characteristics of preparer clientele counts

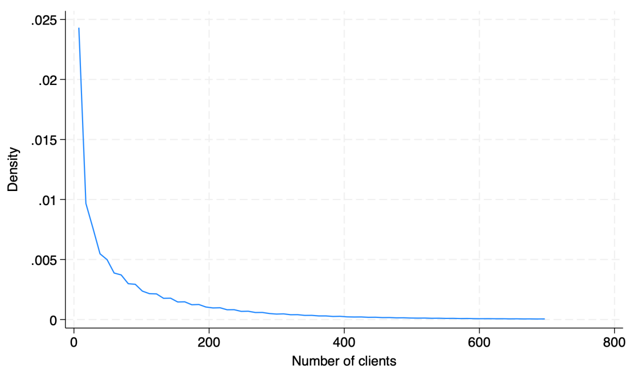

(a) Count of clients per preparer-year

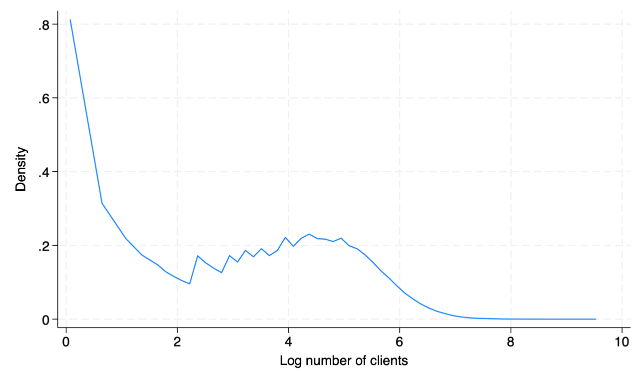

(b) Log count of clients per preparer-year

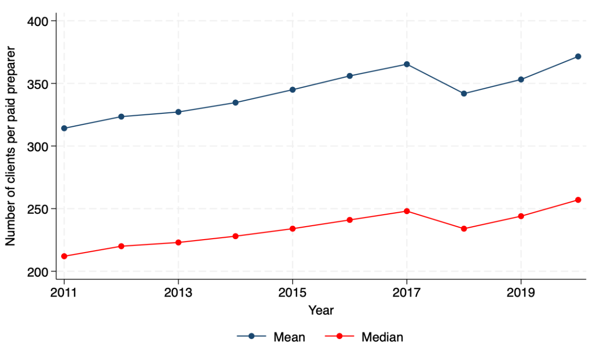

(c) Time series of clients per preparer

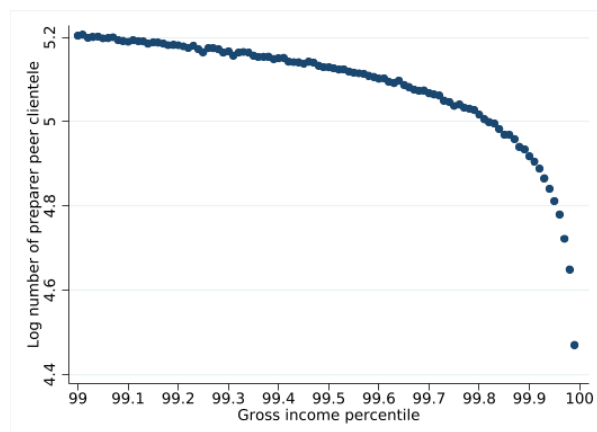

(d) Log number of peer clientele
within the top 1%

Note: These figures show statistics about the number of clientele per preparer-year. Due to confidentiality restrictions, medians are constructed via fuzzy percentiles as means of percentiles 48.8, 48.9, 50, 50.1, and 50.2.

Table A.2: Counterfactual inequality and revenue benchmarks for different allocations of tax preparers

	ETR difference		Proportion tax optimization	
	Unweighted	Weighted	Unweighted	Weighted
Top 1% pre-tax income share	.2427	.2427	.2427	.2427
Benchmark: Top 1% post-tax income share	.2312	.2312	.2312	.2312
All taxpayers:				
Are assigned the mean preparer	.2305	.2305	.2299	.23
Assigned the null-preparer	.2304	.2304	.2288	.2288
Assigned preparer two standard deviations above the mean	.2308	.2307	.2345	.2326
Assigned the mean preparer of the top 1%	.2305	.2304	.2304	.2288
Top 1%:				
Are assigned mean preparer	.2308	.2308	.2302	.2302
Assigned null-preparer	.2297	.2297	.2279	.2279
Assigned preparer two standard deviations above the mean	.2355	.2336	.2395	.2356
Non-users:				
Are assigned mean preparer	.231	.231	.2312	.2312
Assigned preparer two standard deviations above the mean	.2301	.2304	.231	.231
Assigned the mean preparer of the top 1%	.2309	.2312	.2311	.2312

Note: This table displays select estimates for the top 1% post-tax income share and tax revenues raised under different scenarios of the allocation of preparers. Each line corresponds with a counterfactual scenario where we keep the broad income, taxpayer characteristics and year fixed, but replace the tax preparer fixed effect accordingly. The top 1% of earners by broad income is defined on the universe taxpayers prior to the introduction of sample restrictions.

Table A.3: Summary statistics on fees and implied gross/net tax savings

	Mean	SD	p1	p10	p50	p90	p99	N
Fee	362.49	1288.32	22.4	60	210	600	2116.2	101776677
Log fee	5.32	.96	3.1	4.09	5.35	6.4	7.39	101776677
Gross tax savings	1075.57	29728.34	-1010.6	-30.7	245	2108.74	8887.16	101759020
Log gross tax savings	6.44	1.37	2.61	4.79	6.5	8.01	9.1	64510867
Net tax savings	713.25	29628.55	-2001.86	-388.7	-14.52	1636.61	7670.07	101759020
Log net tax savings	6.31	1.48	2.15	4.49	6.39	8	9.15	49663779
Net tax savings share of income	0	1.22	-.02	0	0	.01	.02	101759020

Note: This table displays select summary statistics pertaining to fees and implied gross and net tax savings. Fees are observed on a subset of our data from 2012-2017. Gross tax savings are computed as the tax savings attributable to a preparer, given their estimated ETR difference fixed effect. Net tax savings subtracts the fee from this gross amount. For confidentiality purposes, percentile values correspond with fuzzy percentiles computed as the mean of the five surrounding .1 percentile values (e.g. p50 is computed as the mean of p49.8, p49.9, p50, p50.1, and p50.2).

Figure A.4: Fees metadata: population and selection

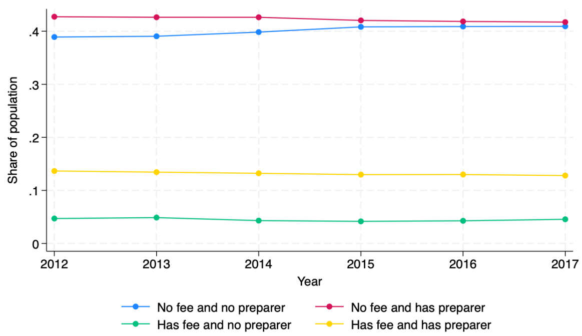

Panel (a): Share of population
by non-missing fee and preparer status

(b) Time series of number of preparers

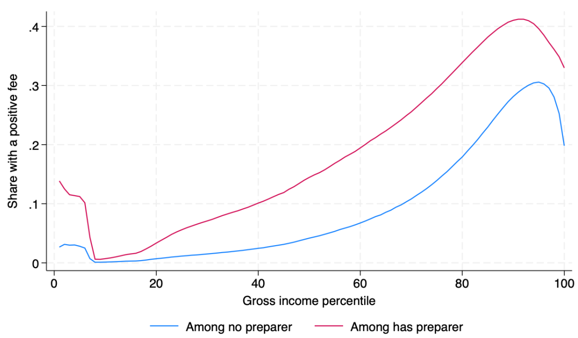

(c) Log income by fees and preparer status

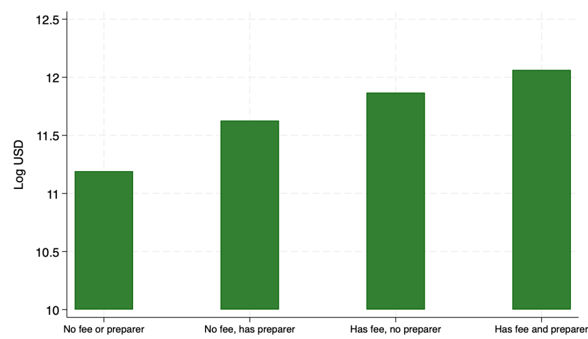

(d) Share of subpopulation
with an observed positive fee

Note: These figures report metadata on fees observation across US taxpayers in our sample from 2012 to 2017. Each year's shares in Panel (a) add to one. Log count in Panel (b) corresponds with the log count of taxpayers in each income percentile bin with a non-missing fee.

Figure A.5: Additional results on fees

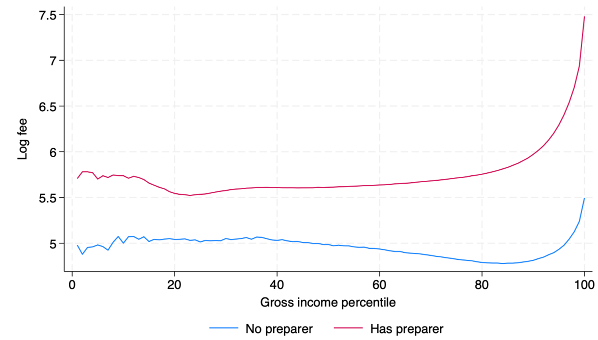

Panel (a): Average log fee by preparer status

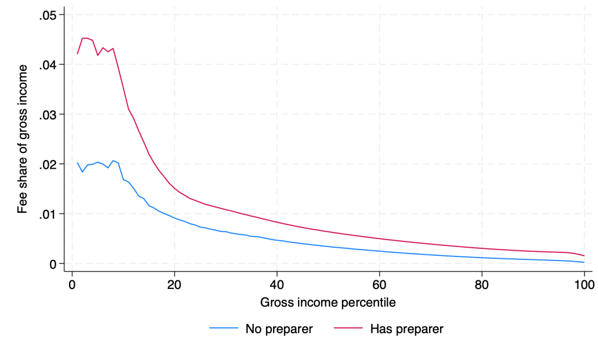

Panel (b) Fee share of broad income by preparer status

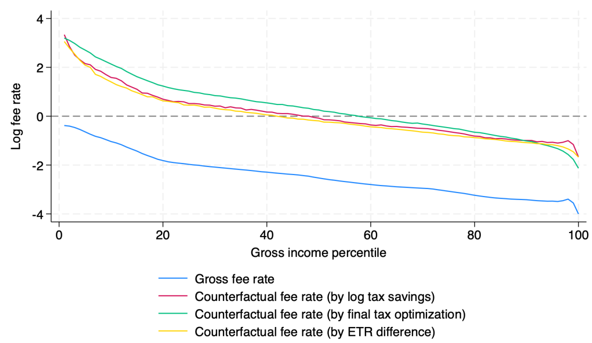

(c) Gross fee rate

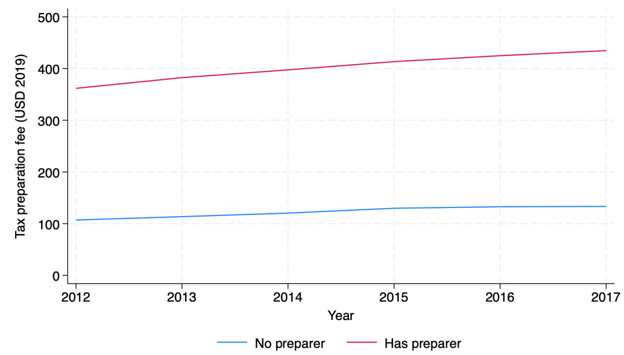

(d) Time series of average fee

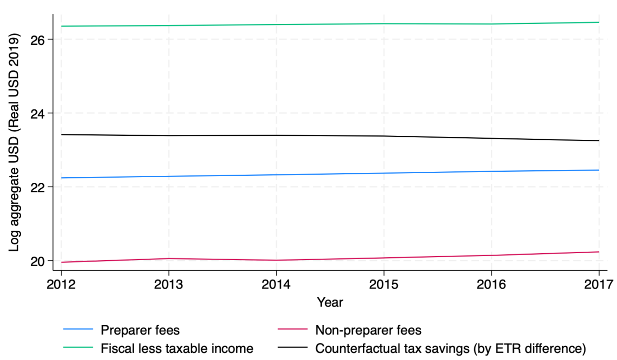

(e) Time series of fees and tax savings

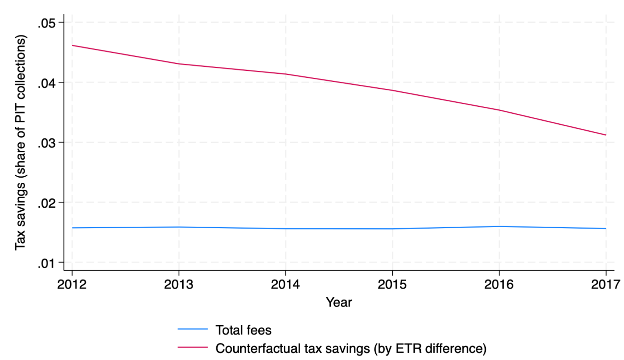

f) Time series of tax savings and fees
Share of gross PIT collections

Note: These figures plot additional results on fees. In Panel (b) fee income shares within each bin are constructed as means across taxpayers as opposed to ratios of aggregated fees and income within bin. Fee rate in Panel (c) corresponds with fee divided by gross tax savings. The counterfactual fee rates estimate the fee rate using savings attributable to the tax preparer relative to the null preparer. Counterfactual tax savings in Panel (f) corresponds with observed preparer fixed effects which are multiplied by taxpayer-year broad income and aggregated and divided by aggregate net personal income tax collections.

Figure A.6: Additional results on fees and preparer quality

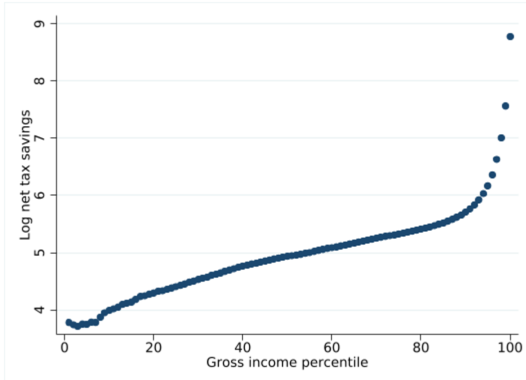

(c) Log net tax savings by income rank

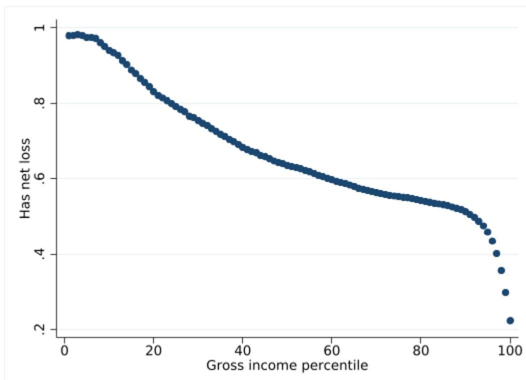

(d) Has a net loss on preparer usage relative to null-preparer

Note: These figures show binned scatter plots of variables of interest along taxpayer gross income percentile. Gross income percentile is computed within-year using broad income. Panel (a) plots log net tax savings, which corresponds with the difference in gross tax savings and fees attributable to the observed tax preparer relative to the null preparer. Panel (b) uses as a dependent variable an indicator for whether preparer tax savings net of fees is negative.

Figure A.7: Operational audit fixed effects: preparers and taxpayers by taxpayer income rank

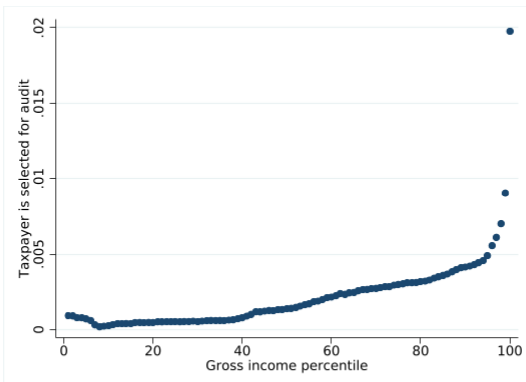

Panel (a): Any operational audit

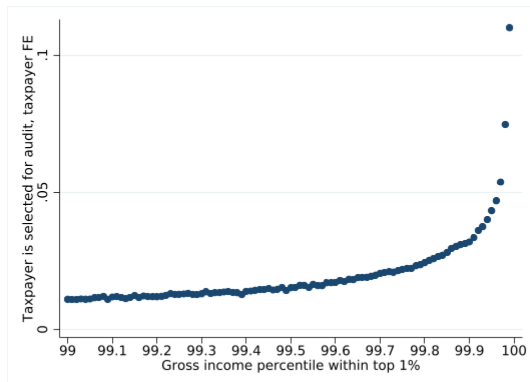

Panel (b): Any operational audit
Top 1% by taxpayer income

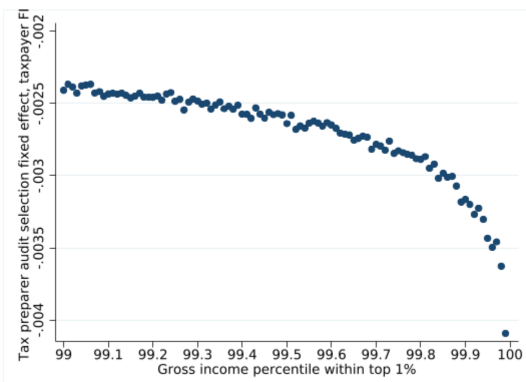

Panel (c): Preparer audit fixed effects
Top 1% by taxpayer income

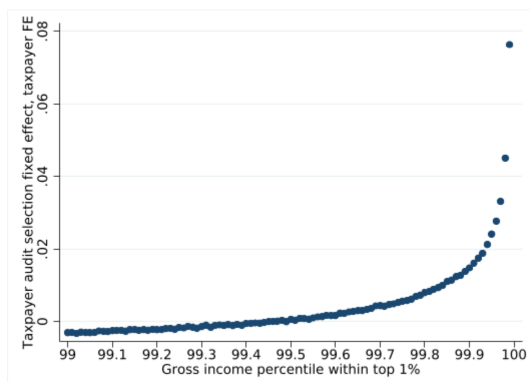

Panel (d): Taxpayer audit fixed effects
Top 1% by taxpayer income

Note: These figures show binned scatter plots of audit selection and audit selection fixed effects along income rank. Panel (a) gives the binned scatter plot for any operational audit conducted along gross income rank; Panel (b) breaks out taxpayers within the top 1%. Panels (c) and (d) plots tax audit selection fixed effects for preparer and taxpayers respectively within the top 1% of earners, elaborating on Figure 7 Panels (a) and (b). Income ranks are computed within-year.

Figure A.8: Audit results: preparer weighting and alternate optimization measures

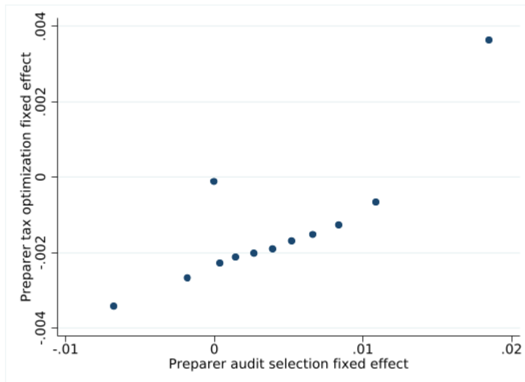

Panel (a): Preparer audit fixed effect versus ETR difference fixed effect: Weighting on $\sqrt{N_j}$

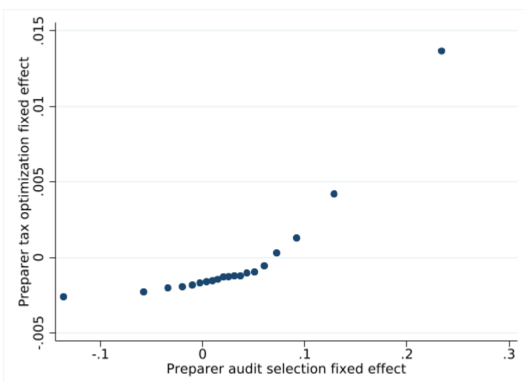

Panel (b): Preparer audit fixed effect versus proportion tax optimization fixed effect: Equal weighting of preparers

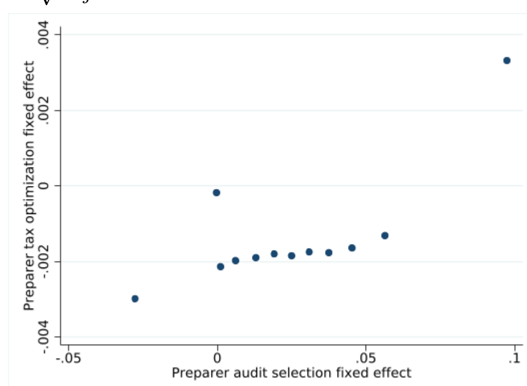

Panel (c): Preparer audit fixed effect versus proportion tax optimization fixed effect: Weighting on $\sqrt{N_j}$

Note: These figures provided robustness to Figure 7 Panel (d). Panel (a) replicates Figure 7 Panel (d) including analytic weights on the number of observations used to compute the fixed effect of preparer j. Panels (b) and (c) use proportion tax optimization fixed effects for tax preparers on the x-axis, equally weighted and using the same weights as in Panel (a) respectively across tax preparers. Panels (a) and (c) feature a discontinuity near the origin reflecting the mass of taxpayers that use the null-preparer. All figures use as observations the unique set of tax preparers with estimated fixed effects.

B Additional information on preparer switches

B.1 Preparer usage

Table B.1: Tabulation of preparer usage from 2011-2020

	(1)	(2)		(3)	(4)		(5)	(6)
Years any preparer used			Different preparers used			Years w/ a given preparer		
	Number	share		Number	share		Number	share
1	27,224,467	18%	1	74,433,507	49%	1	159,950,000	52%
2	17,804,359	12%	2	37,625,222	25%	2	48,098,828	16%
3	14,162,471	9%	3	19,001,652	13%	3	26,097,618	9%
4	12,020,545	8%	4	9,798,187	6%	4	17,334,822	6%
5	10,354,087	7%	5	5,162,913	3%	5	12,522,076	4%
6	9,081,268	6%	6	2,744,145	2%	6	9,554,249	3%
7	8,297,658	5%	7	1,429,969	1%	7	7,569,505	2%
8	7,875,522	5%	8	691,448	0%	8	6,308,714	2%
9	8,825,501	6%	9	277,726	0%	9	5,836,362	2%
10	35,585,926	24%	10	71,254	0%	10	13,637,805	4%

Note: This table tabulates statistics about preparer switches and tenure within our sample of taxpayers from 2011 to 2020. Given by the timespan of our data, the maximum number of unique preparers used or years with a given preparers cannot exceed ten. In columns (3) and (4), the null-preparer is included as a preparer.

Table B.2: Descriptive statistics about preparer switches
Panel (a): Preparer switches over time

		Probability conditional on switch		
Year	Any switch	Preparer-to-preparer	Preparer-to-null	Null-to-preparer
2012	0.190	0.621	0.183	0.197
2013	0.182	0.624	0.187	0.188
2014	0.181	0.633	0.179	0.188
2015	0.177	0.628	0.187	0.185
2016	0.169	0.630	0.178	0.192
2017	0.165	0.633	0.176	0.191
2018	0.164	0.627	0.186	0.187
2019	0.164	0.626	0.202	0.172
2020	0.155	0.623	0.153	0.225

Panel (b): Preparer switches by income tercile

		Probability conditional on switch		
Tercile	Any switch	Preparer-to-preparer	Preparer-to-null	Null-to-preparer
1	0.219	0.558	0.248	0.194
2	0.195	0.605	0.209	0.186
3	0.145	0.658	0.147	0.195

Note: These tables give tabulations of preparer switch probabilities over time and within income tercile. A switch is defined for a taxpayer i in year t as two consecutive years in which a fixed individual is observed with two different: $J(i, t-1) \neq J(i, t)$. As such, switches cannot be computed in 2011, the first year of our data. Switches can be partitioned between preparer-to-preparer, preparer-to-null and null-to-preparer switches, so that conditional probabilities sum to one. Income terciles are computed within-year.

Table B.3: Characteristics of preparer usage

	Mean	SD	p5	p25	p50	p75	p95	N
Preparer usage								
Number of years in the panel	5.604	3.294	1	2	6	9	10	1.710e+08
Number of years with any preparer	3	3.417	0	0	1	5	10	1.710e+08
Fraction of years with any preparer	0.516	0.449	0	0	0.500	1	1	1.710e+08
Ever self-prepared	0.615	0.487	0	0	1	1	1	1.710e+08
Ever used a preparer	0.643	0.479	0	0	1	1	1	1.710e+08
Number unique preparers	1.740	1.150	1	1	1	2	4	1.710e+08
Number unique preparers (non-null)	1.126	1.246	0	0	1	2	4	1.710e+08
Raw years per preparer	3.873	2.899	1	1.500	3	5	10	1.710e+08
Years per preparer (conditional on usage)	3.104	2.645	1	1	2	4	10	1.100e+08
Switch probability	0.217	0.299	0	0	0	0.333	1	1.450e+08
Years with outgoing preparer	2.357	1.802	1	1	1.750	3	7	7.170e+07
Years with incoming preparer	2.490	1.846	1	1	2	3	7	7.170e+07
Conditional on switch:								
Switch probability 1-1	0.540	0.452	0	0	0.600	1	1	7.170e+07
Switch probability 0-1	0.229	0.346	0	0	0	0.500	1	7.170e+07
Switch probability 1-0	0.231	0.346	0	0	0	0.500	1	7.170e+07
Years with outgoing preparer (1-1 switch)	2.192	1.741	1	1	1.500	2.667	6	4.530e+07
Years with incoming preparer (1-1 switch)	2.235	1.738	1	1	1.500	3	6	4.530e+07
Years with outgoing preparer (0-1 switch)	2.667	2.020	1	1	2	4	7	2.690e+07
Years with incoming preparer (0-1 switch)	2.028	1.695	1	1	1	2	6	2.690e+07
Years with outgoing preparer (1-0 switch)	1.767	1.405	1	1	1	2	5	2.710e+07
Years with incoming preparer (1-0 switch)	2.834	2.071	1	1	2	4	7	2.710e+07

Note: This table displays characteristics pertaining to taxpreparer usage among our baseline sample prior to restricting to the leave-one-out largest connected set of taxpayers and tax preparers. In the panel below "conditional on switch:", the transitions indicated "x-y" correspond with preparer or null-preparer movements, where 1 corresponds with paid preparers and 0 corresponds with the null-preparer. For confidentiality purposes, percentile values correspond with fuzzy percentiles computed as the mean of the five surrounding .1 percentile values (e.g. p50 is computed as the mean of p49.8, p49.9, p50, p50.1, and p50.2).

Figure B.1: Preparer switches over time

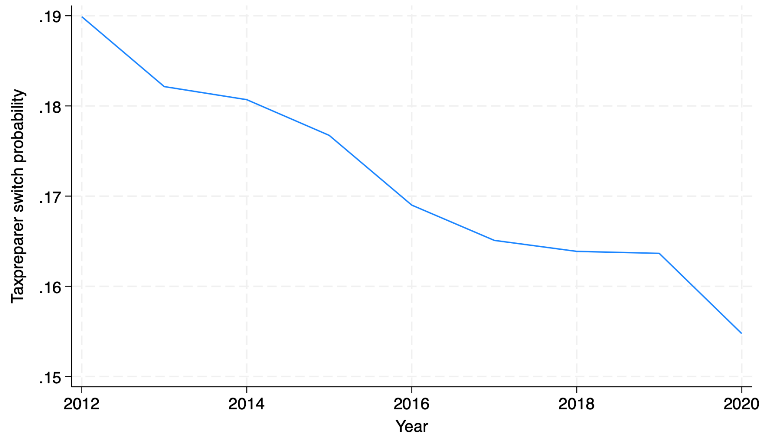

(a) Switch probability over time

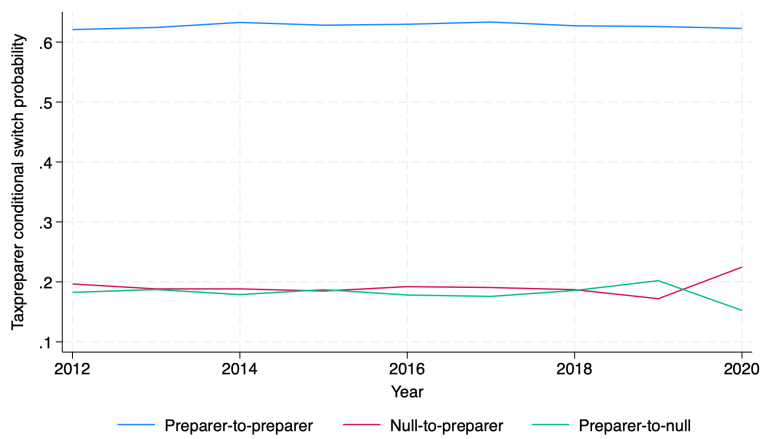

(b) Conditional probability of each switch type

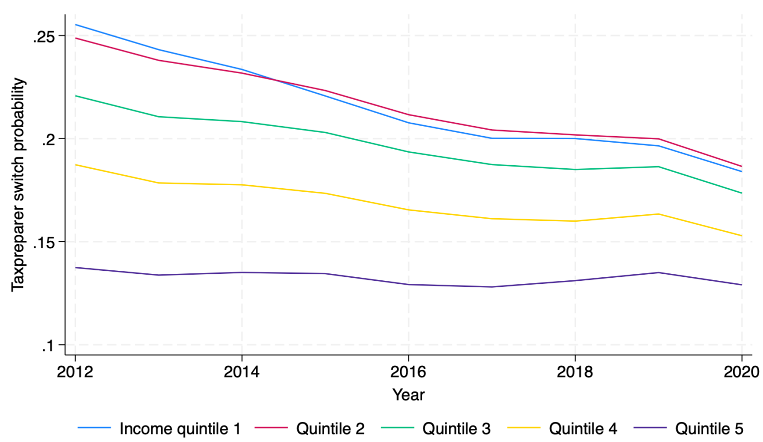

(c) Switch probability by income quintile

Note: these figures tabulate preparer switch characteristics on our leave-out-out largest connected set sample of taxpayers. A switch is defined as any observation where of fixed taxpayer reports two different tax preparer IDs in two consecutive years, for which reason these time series only begin in 2012.

Figure B.2: Evolution of taxpayer outcomes around preparer moves:
By incumbent and destination preparer leave-out mean ETR difference
Non-residualized tax optimization

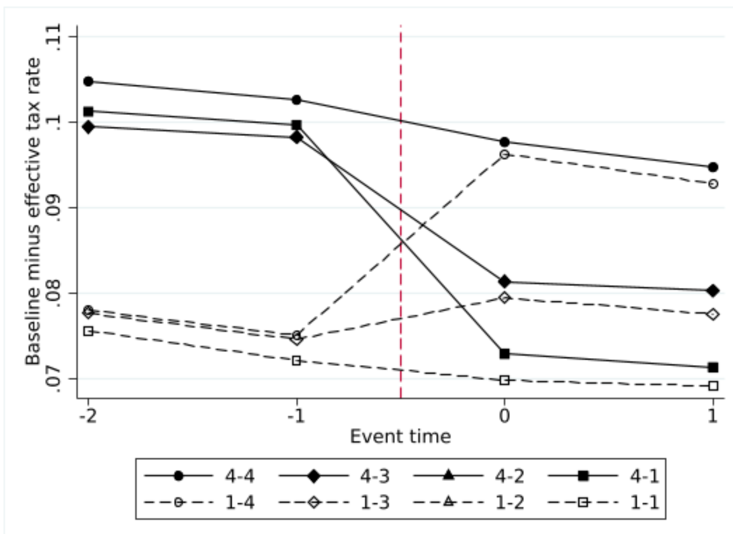

ETR Difference

Note: Note: This figure shows mean levels of ETR difference for taxpayers that switch tax preparers. We consider the set of taxpayer-preparer switches in which the incumbent and incoming match lasted at least two years. Each preparer is classified into quartiles based on the mean wage of preparer peer-clientele in the same year. Within each move-type, we plot the average ETR difference by event-time. The dependent variable is constructed as the percentage point difference between ETR based on statutory tax imputed from broad income and final realized ETR.

Figure B.3: Placebo test: evolution of taxpayer broad earnings around preparer moves: By incumbent and destination preparer leave-out mean ETR difference

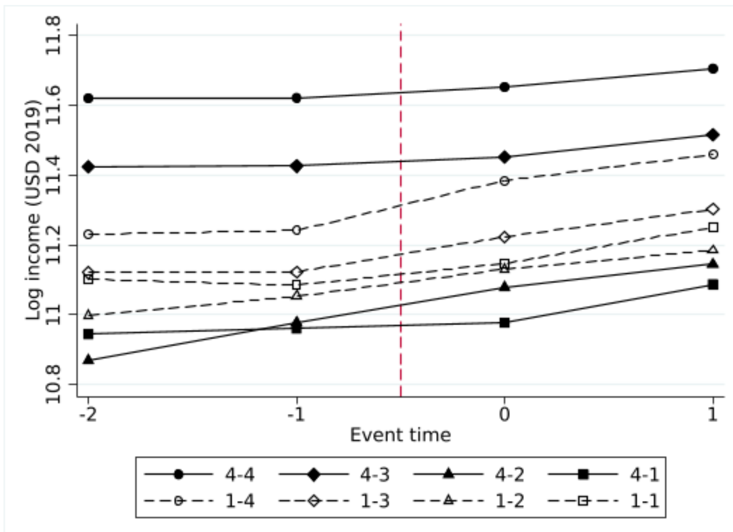

Panel (b): Placebo test: Gross income

Note: Note: This figure shows mean levels log broad income for taxpayers that switch tax preparers. We consider the set of taxpayer-preparer switches in which the incumbent and incoming match lasted at least two years. Each preparer is classified into quartiles based on the mean value of ETR difference of preparer peer-clientele in the same year. Within each move-type, we plot the average ETR difference by event-time. The dependent variable is constructed as the percentage point difference between ETR based on statutory tax imputed from broad income and final realized ETR

C Replication of baseline results for alternative tax optimization measures

C.1 Description of the alternative measures

We also define an alternate measure tax optimization as:

$$\text{Proportional tax optimization} := \frac{T^b - T^f}{T^b} = 1 - \frac{T^f}{T^b}, \tag{7}$$

This alternate measure corresponds to the proportion difference in final taxes paid relative to as imputed statutorily from broad income using our mapping. We define this measure so that, syntactically, greater positive values correspond with greater levels of tax savings. E.g. a value of 0.4 implies that final taxes paid were 40% lower than the amount imputed from broad income. This measure also has the tractable quality that it is the proportion difference (as opposed to levels difference) in ETR between final taxes paid relative to imputed taxes from broad income: Proportion tax optimization $= \frac{\text{ETR difference}}{T^b/B} = \frac{\frac{T^b}{B} - \frac{T^f}{B}}{\frac{T^b}{B}}$. There are important tradeoffs with these two measures. We view ETR difference as likely more intuitive and easy to work with for back-of-the-envelope calculations. Proportion tax optimization is more considerate of proportional differences in ETR at different parts of the income distribution, but performs poorly at very low levels of tax liability. ETR difference also performs poorly for tax returns with near-zero levels of broad income; however, we remove such individuals from our sample, as discussed in Section 3.

We consider other measures as well. Two of note here are 1) "Log tax savings", which we define as $log(T^f - T^b)$, and 2) "Final base optimization", which relies on an inverse function concept to map values of taxes to amounts of income tax base. Log tax savings models tax savings, similar to wages, as log-normally distributed (Mincer, 1958). The measure benefits from tractable properties of the logarithm (e.g. approximating proportion differences at values near zero), but its performance begins to deteriorate when there are differences between T^b and T^f in at least approximately one order of magnitude (where the proportion approximation of the logarithm begins to fail), which likely occurs at differentially along the income distribution.

The last other measure is “proportion base optimization” which operates similar to proportion tax optimization, but instead of taxes in the numerator, gross income and imputed gross income (from tax payments) serve as the objects of interest. We can define “Proportion base optimization := $\frac{B-T^{-1}(T^f)}{B} = 1 - \frac{T^{-1}(T^f)}{B}$. One reason why analyzing “base optimization” might be more informative than “tax optimization” in quantifying an amount of tax strategy is that due to the nonlinearity of the income tax schedule, an identical application of deductions and credits could result in different tax obligations starting from different points of broad income. Therefore, analyzing the amount of deductions and credits applied themselves may be more informative of “effort” in tax strategy than the final result in taxes due. On the other hand, from a money-metric perspective, final taxes paid are arguably more relevant for taxpayer consumption and objective maximization. However, this measure relies on a forced bijection between income and taxes, which is potentially problematic. This inverse mapping T^{-1} keeps fixed filing characteristics and treats all income as taxed as standard ordinary wage income. Because tax base contains several different types of income that receive different tax treatment, such as for example, ordinary wage income or qualified dividend income, the mapping $T(B;\Theta)$ is not truly bijective. To be able to construct our mapping, we treat all income as wage income. Of course, the tax rate on wage can be higher or lower than tax on capital across the income distribution, but in practice, ordinary wage income sees a higher statutory tax rate than capital income or other types of income, so we argue wage income constitutes a suitable baseline.

Lastly, one may be interested in distinguishing between credits and deductions in the optimization process. This kind of comparison yields several complications due to the nonlinearity of the income tax function: deductions and credits are expressed in fundamentally different units, and require conversion for comparison. One can map deductions into tax-value money metric using the difference in the imputed income taxes at the point of application of the deduction $D_T = T(B) - T(B - D)$. Alternatively, one can map credits into deductions-value or tax base equivalent using the inverse mapping from taxes to income and a similar approach: $C_B = T^{-1}(T^f + C) - T^{-1}(T^f)$. However, one also needs to note the refundability or non-refundability of credits. In the case of non-refundable credits, the base-equivalent value of credits cannot exceed the base-equivalent value of tax obligation prior to the application of credits, whereas for refundable credits, the base-equivalent value is unbounded.

C.2 Results

Figure C.2 uses proportion tax optimization as the main dependent variable. The figure shows that tax optimization is progressive, decreasing monotonically in income rank; this aligns with the main findings in the respect that lower earning taxpayers owe far less in taxes, so that proportionately, an identical reduction in percentage points ETR represents a larger proportional reduction. Panel (a) also documents a regressive turn in proportional tax optimization at the very top of the income distribution. Panels (b) - (d) carry an identical interpretation to their respective analogue in Figure 2.

Figure C.1: Distribution of proportion tax optimization

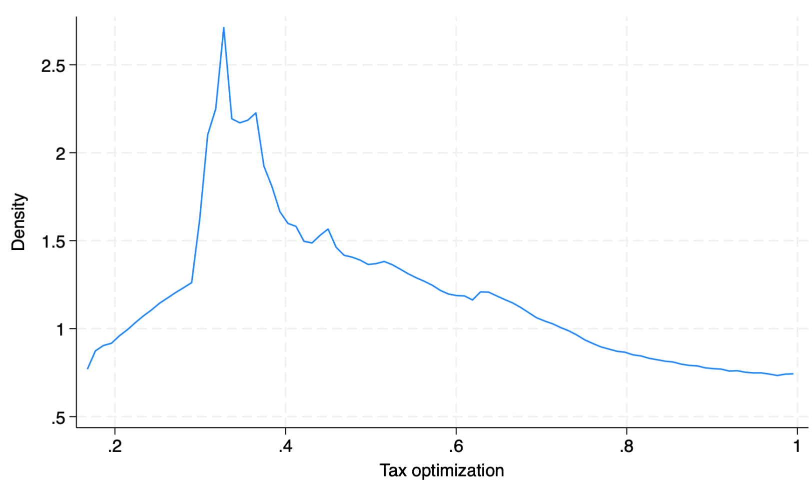

Note: This figure displays a histogram the variable proportion tax optimization. Each observation corresponds with a taxpayer-year.

Figure C.2: Characteristics of tax optimization and preparer usage
Proportion tax optimization

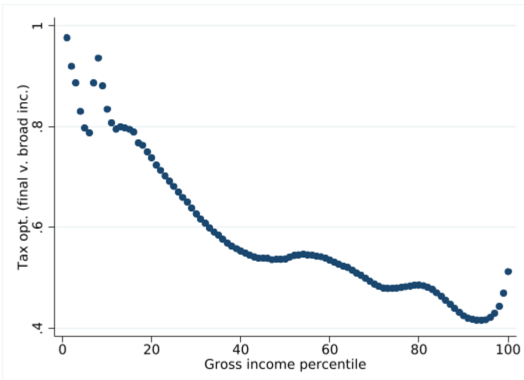

Panel (a): ETR difference along
the broad income distribution

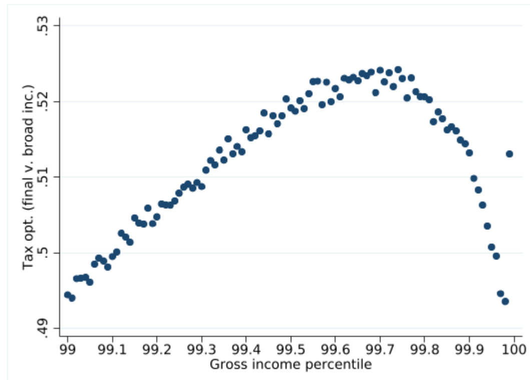

Panel (b) ETR difference within
the top 1% of broad income

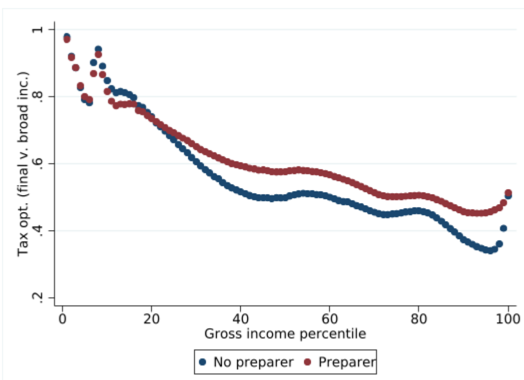

Panel (c): Proportion tax optimization along
the broad income distribution:
Preparer users v. non-users

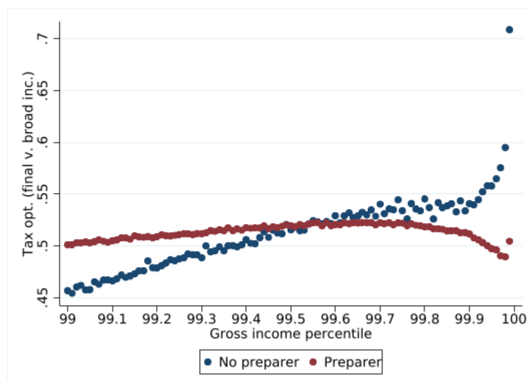

Panel (d) Proportion tax optimization
the top 1% of broad income:
Preparer users v. non-users

Note: These figures display binned scatter plots of various characteristics of paid preparer usage along gross income percentile. Gross income percentiles are constructed as within-year percentile ranks of tax-filing units, where gross income represents income across different reported sources within the income tax declaration prior to the application of credits and deductions. Panels (c) and (d) plots the binned scatter plots in Panel (a) and (b) respectively, stratifying by paid preparer status.

Figure C.3: Evolution of taxpayer outcomes around preparer moves:
By incumbent and destination preparer leave-out mean proportion tax optimization

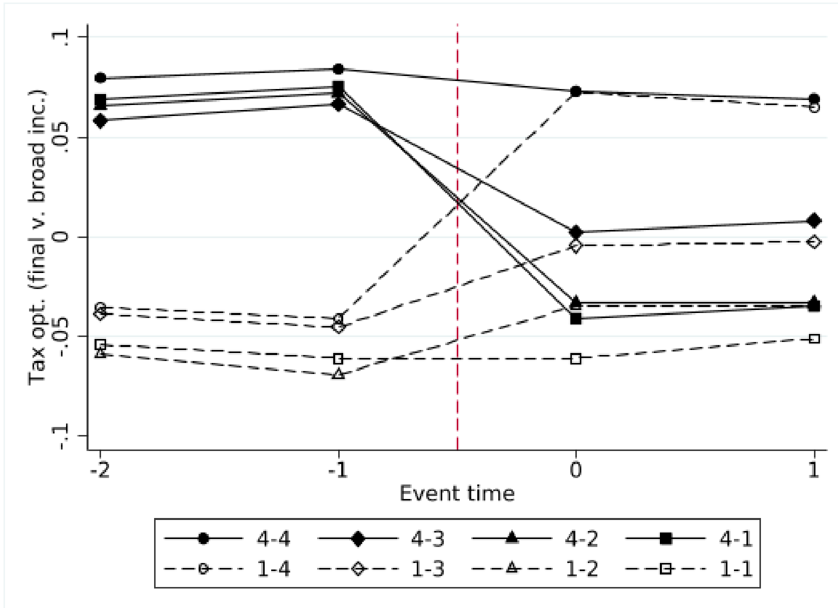

(a) Proportion tax optimization

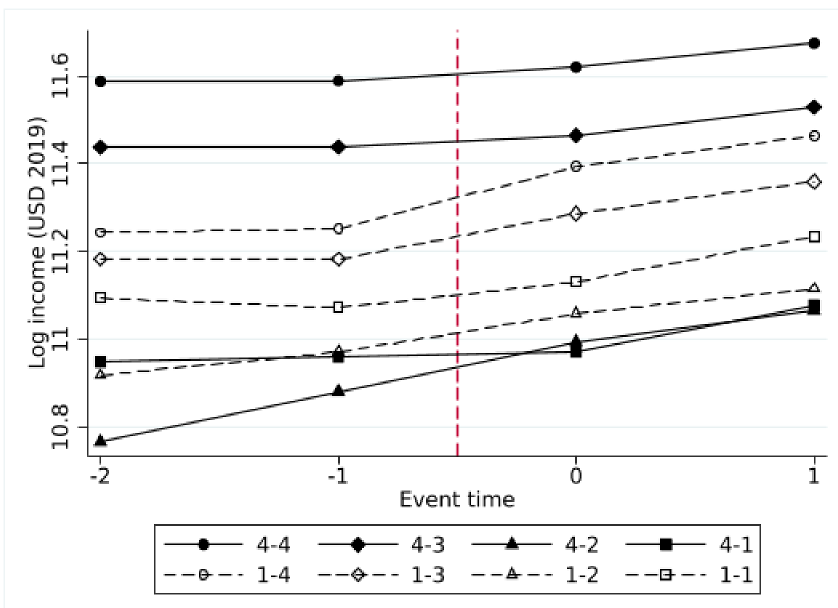

(b) Placebo test: Gross income

Note: These figures show mean levels of proportion tax optimization and log broad income, respectively, for taxpayers that switch tax preparers. We consider the set of taxpayer-preparer switches in which the incumbent and incoming match lasted at least two years. Each preparer is classified into quartiles based on the mean value of proportion tax optimization difference of preparer peer-clientele in the same year. Within each move-type, we plot the average of the dependent variable by event-time. Proportion tax optimization is constructed as the proportion difference between ETR based on statutory tax imputed from broad income and final realized ETR, which is residualized on year and state fixed effects and log income.

Table C.1: Variance Components of Tax Optimization (appendix)

	All	No null preparers	By Income Low	Medium	High
	(1)	(2)	(3)	(4)	(5)
(i) Tax opt final					
STD preparers (ψ)	0.026	0.032	0.022	0.026	0.024
STD taxpayers (α)	0.157	0.172	0.135	0.158	0.170
$corr(\psi, \alpha)$	0.096	0.019	0.018	0.060	0.090
# of preparers	1,267,867	1,267,866	1,068,091	1,049,060	947,050
# of taxpayers	171,237,396	111,417,835	78,770,676	51,150,655	41,316,065

Note: This table displays standard deviations and correlations between different fixed effect objects from our two-way fixed effects estimation procedure. The dependent variable in all specifications is proportion tax optimization. Column (1) displays estimates from our main specification. Column (2) contains estimates from an alternate specification that omits taxpayers without a paid tax preparer. Columns (3) - (5) stratify our estimation procedure by within-year broad income tercile. Each specification indicates the final number of unique tax preparers and tax payers that constitute the leave-one-out largest connected set.

D Additional results about preparer usage

Table D.1: Summary Statistics on preparer's Clientele by Professional Category

	Attorney	CAA	CPA	EA	SRTP
Number of annual clientele	2.194	3.855	2.858	3.569	3.819
	(2.299)	(4.081)	(3.117)	(3.537)	(4.315)
AGI	104241	41822	119271	75310	55548
	(108435.1)	(38091.86)	(106934.6)	(66638.44)	(48718.19)
Tax obligation	16013	2413	19050	9091	4837
	(25056.36)	(7226.771)	(25195.08)	(14491.44)	(9929.234)
Average tax rate	0.086	0.048	0.092	0.075	0.059
	(.072)	(.041)	(.065)	(.053)	(.048)
Tax opt. (total)	0.499	0.624	0.492	0.514	0.584
	(.437)	(.306)	(.371)	(.345)	(.328)
Base opt. (total)	0.471	0.594	0.456	0.479	0.549
	(.387)	(.298)	(.331)	(.313)	(.309)
Ded. + cred. inc. share	0.399	0.697	0.396	0.397	0.487
	(.953)	(1.137)	(.875)	(.748)	(.844)
Distinct preparers	7525	3310	131387	47642	23390
Preparer-years	37524	20100	824691	302201	134648

Note: This table disaggregates paid preparer clientele characteristics by paid preparer type. CAA refers to "Certified Acceptance Agent". CPA refers to "Certified Public Accountant"; EA refers to "enrolled agents". SRTP refers to "Supervised Registered Tax Preparer". Standard errors are clustered on the taxpayer level.

Table D.2: Taxpayers × Preparer Usage

	Mean	SD	p5	p25	p50	p75	p95	N
Preparer usage								
Number of years in the panel	5.604	3.294	1	2	6	9	10	1.710e+08
Number of years with any preparer	3	3.417	0	0	1	5	10	1.710e+08
Fraction of years with any preparer	0.516	0.449	0	0	0.500	1	1	1.710e+08
Ever self-prepared	0.615	0.487	0	0	1	1	1	1.710e+08
Ever used a preparer	0.643	0.479	0	0	1	1	1	1.710e+08
Number unique preparers	1.740	1.150	1	1	1	2	4	1.710e+08
Number unique preparers (non-null)	1.126	1.246	0	0	1	2	4	1.710e+08
Raw years per preparer	3.873	2.899	1	1.500	3	5	10	1.710e+08
Years per preparer (conditional on usage)	3.104	2.645	1	1	2	4	10	1.100e+08
Switch probability	0.217	0.299	0	0	0	0.333	1	1.450e+08
Years with outgoing preparer	2.357	1.802	1	1	1.750	3	7	7.170e+07
Years with incoming preparer	2.490	1.846	1	1	2	3	7	7.170e+07
Conditional on switch:								
Switch probability 1-1	0.540	0.452	0	0	0.600	1	1	7.170e+07
Switch probability 0-1	0.229	0.346	0	0	0	0.500	1	7.170e+07
Switch probability 1-0	0.231	0.346	0	0	0	0.500	1	7.170e+07
Years with outgoing preparer (1-1 switch)	2.192	1.741	1	1	1.500	2.667	6	4.530e+07
Years with incoming preparer (1-1 switch)	2.235	1.738	1	1	1.500	3	6	4.530e+07
Years with outgoing preparer (0-1 switch)	2.667	2.020	1	1	2	4	7	2.690e+07
Years with incoming preparer (0-1 switch)	2.028	1.695	1	1	1	2	6	2.690e+07
Years with outgoing preparer (1-0 switch)	1.767	1.405	1	1	1	2	5	2.710e+07
Years with incoming preparer (1-0 switch)	2.834	2.071	1	1	2	4	7	2.710e+07

Note: This table presents descriptive information on preparers usage by taxpayers from 2011 to 2019. In addition to the mean (reported in the first column), the table reports distributional information across several percentiles. For confidentiality purposes, percentile values correspond with fuzzy percentiles computed as the mean of the five surrounding .1 percentile values (e.g. p50 is computed as the mean of p49.8, p49.9, p50, p50.1, and p50.2).