A Welfare Analysis of Tax Audits Across the Income Distribution

William C. Boning, U.S. Department of the Treasury
Nathaniel Hendren, MIT and Policy Impacts
Ben Sprung-Keyser, University of Pennsylvania and Policy Impacts
Ellen Stuart, The University of Sydney and Policy Impacts

August 2024

Online Appendix

Contents

A Appendix Figures **3**

B Data Appendix **22**
B.1 Audit-Level Enforcement Data: Revenue and Direct Labor Cost Estimates 22
B.2 Accounting Data: Non-Direct Labor Cost Estimates 24

C Correspondence Audits **32**

D Relation to Existing Estimates **32**

E MVPF of Tax Evasion **34**
E.1 Setup 34
E.2 Willingness to Pay for Expanded Audits 35
E.3 Government Revenue and the MVPF of Expanded Audits 36
E.4 Robustness and Extensions 38
E.4.1 Risk Aversion 39
E.4.2 General Deterrence 40
E.4.3 Endogenous Audit Probability 41
E.5 Comparing the MVPF of Audits to the MVPF of Tax Rate Changes 45
E.6 Heterogeneous Welfare Weights and Comparisons Across Policies 46
E.7 Audit Rates versus Audit Penalties 48

A Appendix Figures

APPENDIX FIGURE A.I:

Overview of the Audit Process

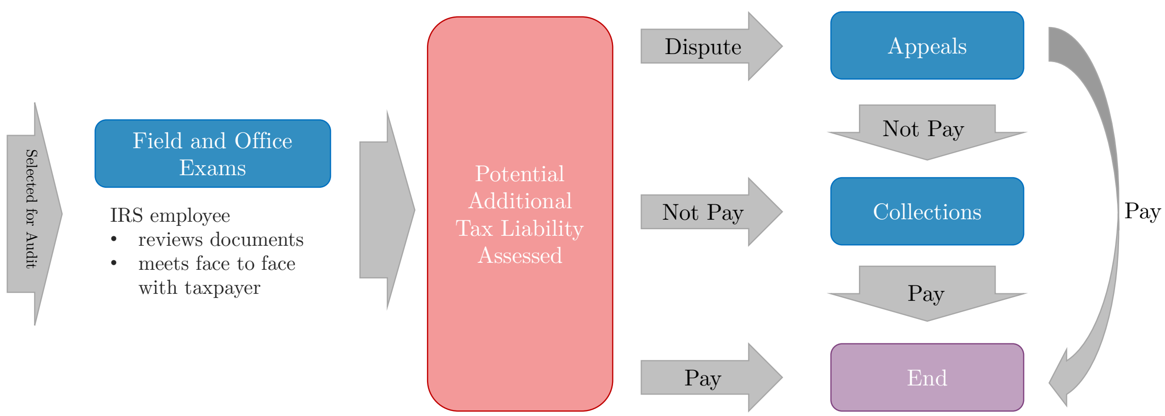

Notes: This flow chart provides an overview of the audit process beginning when a tax return is selected for a field or office exam. The exam begins when an IRS employee reviews the taxpayer's relevant documents and meets face-to-face with the taxpayer to determine any adjustments to tax liability. If there is no audit adjustment then the enforcement process concludes. If the taxpayer agrees with the audit adjustment then the change in tax liability is assessed, while if the taxpayer disagrees the case is heard by the IRS' Independent Office of Appeals and may be further appealed to tax court for a final determination and assessment. If the taxpayer does not pay any additional assessed liability, the case is sent to collections. In practice, few exams are appealed or sent to collections. Our estimates include the taxes, penalties, and interest collected and costs accrued at the exam, appeals, and collections stages of the audit process.

APPENDIX FIGURE A.II:

Average Costs and Revenue Raised per Correspondence Audit

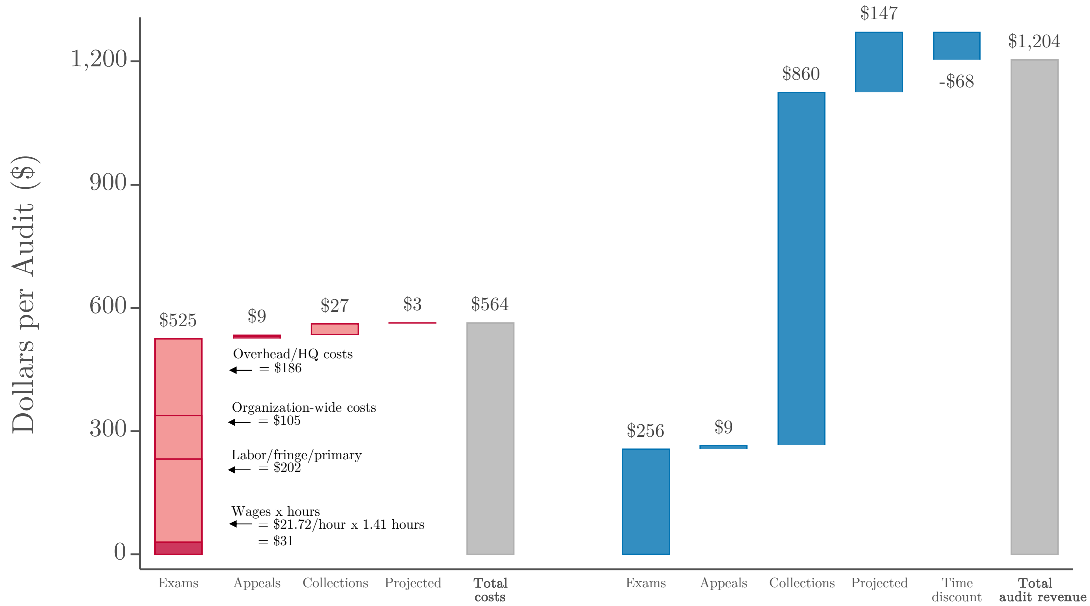

Notes: This figure presents the average total costs and revenue raised per correspondence audit of a tax return filed for tax years 2010–2014. Total wage costs (auditors' wages times hours spent on exam) are shown in dark red and additional costs are shown stacked on top in lighter red. Additional costs include labor/fringe/primary, organization-wide, and overhead/HQ costs. Together, these additional costs are 16.07, 0.57, and 31.83 times total wage costs at the exam, appeals, and collections stages respectively using average multiplier values from tax years 2011–2015. Revenue raised at each stage of the audit process is shown in blue and includes revenue raised from additional tax liability, penalties and interest. Average costs and revenues include projected costs incurred and revenue collected after the observed 7–11 year post-audit sample window. Revenues are discounted using a 3% discount rate because revenues lag costs by about a year on average. In particular, we use data from the 2003 tax year to separately discount the revenues raised and costs accrued each year post-audit back to the tax year. We then use the ratio of the discounted series (net present value of revenues over costs) to adjust revenues downwards to align the two paths.

APPENDIX FIGURE A.III:

Correspondence Audits Across the Income Distribution

A. Audit Rates

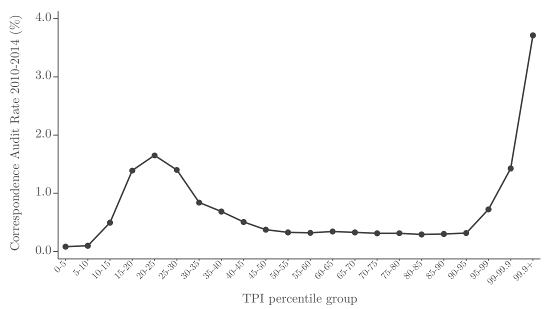

B. Wage Costs

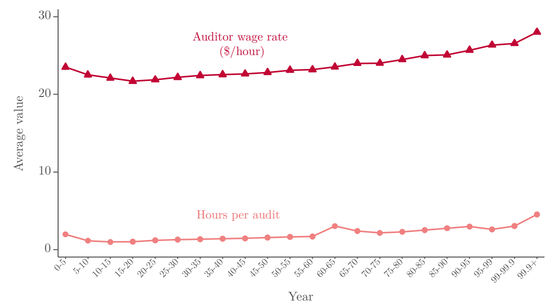

C. Audit Revenue and Total costs

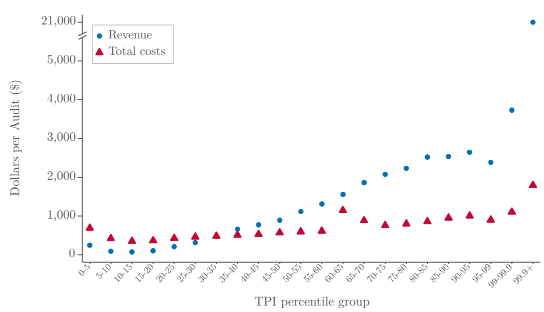

D. Audit revenue over Total Costs

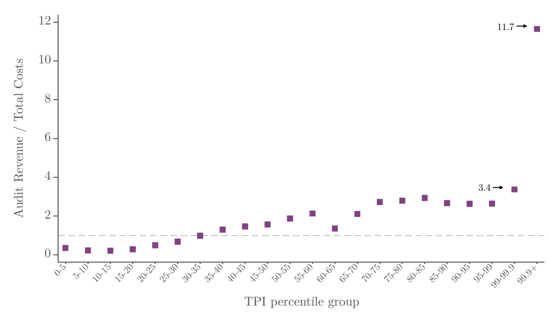

Notes: This figure presents the average audit rate, costs and revenues for correspondence audits of tax returns filed for tax years 2010–2014 by the taxpayer's total positive income (TPI). Panel A presents the average audit rate for correspondence audits. Panel B shows each component of labor costs (auditors' wages and hours worked per audit) per correspondence audit. Panel C presents the average total costs and revenue raised per audit, and Panel D shows the ratio of the average revenue and costs per audit by TPI. The x-axis groups TPI into bins of five percentiles and splits out the top bin into the 95–99th and 99–99.9th percentiles and the top 0.1%. Total costs are the sum of labor costs (auditors' wages times hours spent on exam) and additional costs (labor/fringe/primary, organization-wide, and overhead/HQ costs) which are allocated in proportion to direct labor costs. Total revenue is the sum of additional tax liability, penalties and interest collected. Average costs and revenues include projected costs incurred and revenue collected after the observed 7–11 year post-audit sample window. Revenues are discounted using a 3% discount rate because revenues lag costs by about a year on average. In particular, we use data from the 2003 tax year to separately discount the revenues raised and costs accrued each year post-audit back to the tax year. We then use the ratio of the discounted series (net present value of revenues over costs) to adjust revenues downwards to align the two paths.

APPENDIX FIGURE A.IV:

Average Costs, Revenue and Revenue over Costs per In-Person Audit with Alternative Non-Direct Labor Cost Allocations, by Income Group

A. Average Costs and Revenue

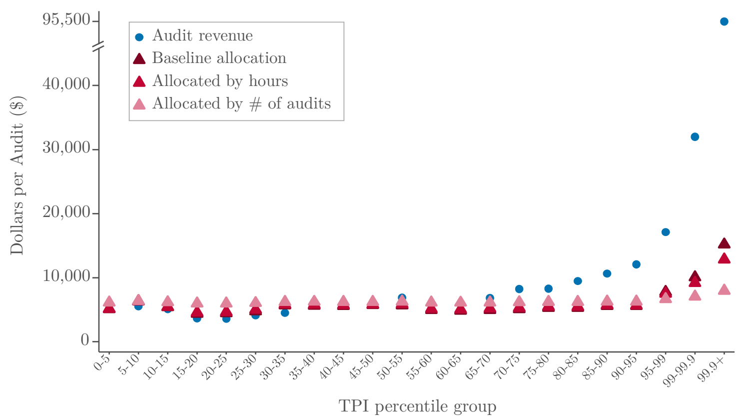

B. Average Revenue over Costs

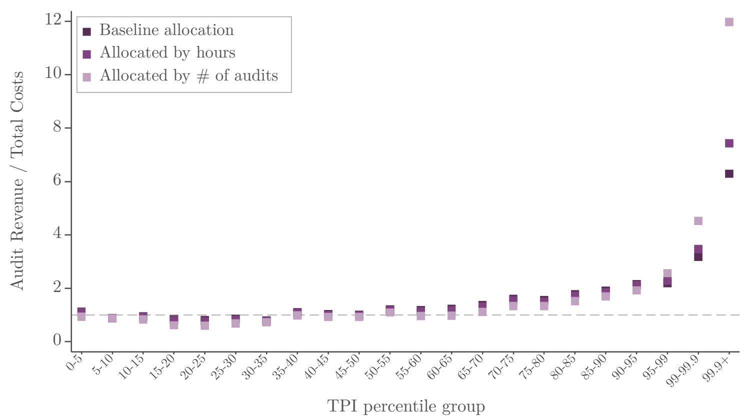

Notes: This figure presents the average total costs accrued and revenue raised per in-person audit (Panel A) and the ratio of average revenue and costs (Panel B) by the taxpayer's total positive income (TPI) using different methods to allocate overhead costs (non-direct labor-related costs, organization-wide costs, and general overhead costs) across the income distribution. The baseline method shown in dark purple allocates overhead costs in proportion to direct audit wage costs. The second method shown in the mid shade of purple allocates costs in proportion to labor hours rather than total labor costs. The third method shown in light purple allocates overhead costs equally per audit.

APPENDIX FIGURE A.V:

Cumulative Share of Revenue Collected and Labor Costs Accrued, by Years Post-Tax Year and Income

A. Audit Revenues

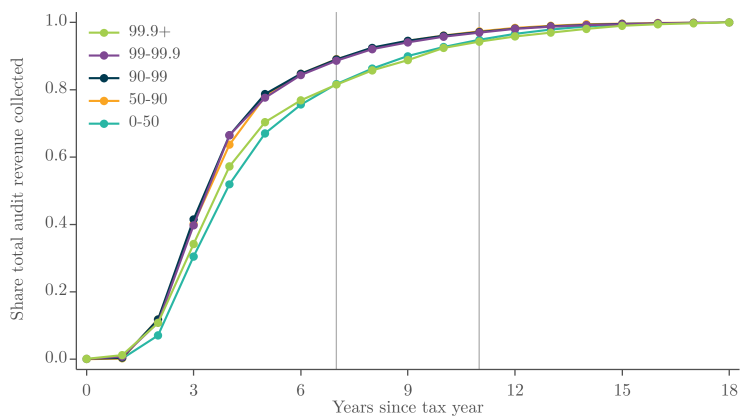

B. Labor Costs

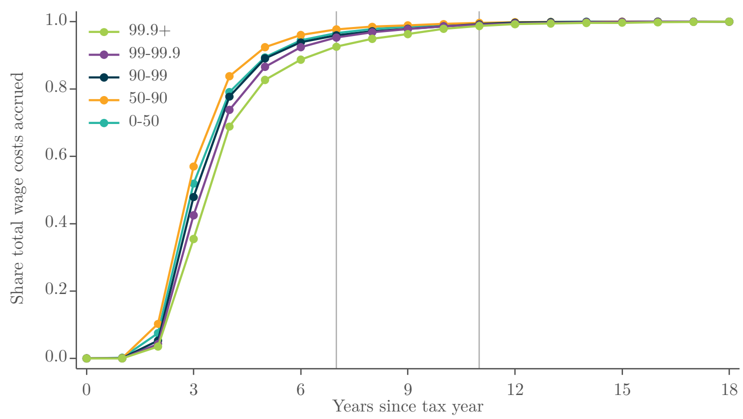

Notes: This figure presents the trajectory of revenues raised (Panel A) and costs accrued (Panel B) following audits of tax year 2003 returns. The 2003 tax year lies before our primary sample window, but shows 18 years of follow-up data. The y-axis shows the cumulative share of revenues collected and labor costs accrued relative to the total values 18 years post-tax year. The gray vertical lines indicate the 7–11 year windows observed after the tax year 2010 to 2014 returns in our primary sample were filed.

APPENDIX FIGURE A.VI:

Average Costs, Revenue and Revenue over Costs per In-Person Audit, by Year

A. Average Costs and Revenue

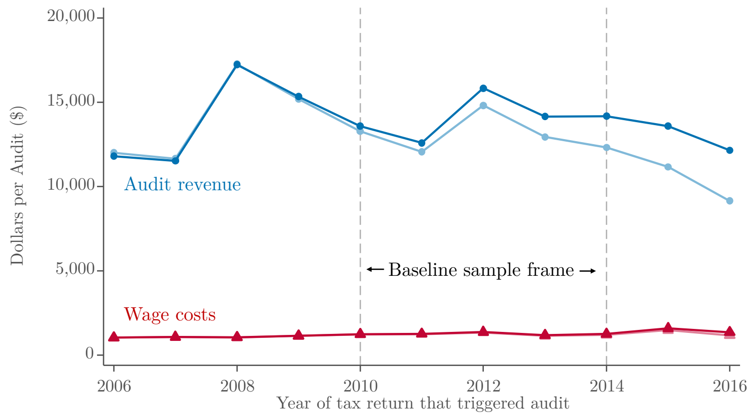

B. Average Revenue over Costs

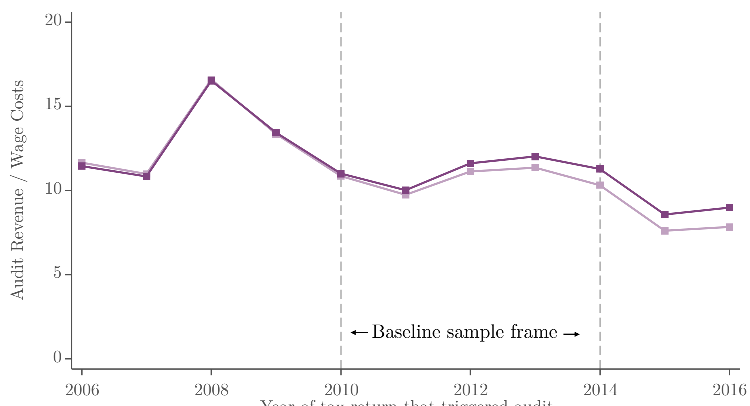

Notes: Panel A presents the average total costs and revenue raised per in-person audit by the tax year for which a return was filed. Panel B shows the ratio of the average revenue and costs per audit by tax year. The vertical gray dashed lines indicate our primary sample window. Total costs are the sum of labor costs (auditors' wages times hours spent on exam) and additional costs (labor/fringe/primary, organization-wide, and overhead/HQ costs). Total revenue is the sum of additional tax liability, penalties and interest collected. Average costs and revenues include projections of revenue collected and costs accrued outside the observed post-audit sample window for each tax year. The series plotted in the lighter shades of blue, red, and purple show the average values of revenues, costs, and revenues over costs without this projection adjustment. Revenues are discounted using a 3% discount rate since average revenues lag average costs by approximately one year. In particular, we use data from the 2003 tax year to separately discount the revenues raised and costs accrued each year post-audit back to the tax year. We then use the ratio of the discounted series (net present value of revenues over costs) to adjust revenues downwards to align the two paths.

APPENDIX FIGURE A.VII:

Average In-Person Audit Revenue Over Total Costs Using Imputed Post-Audit TPI

A. TPI Imputed by Applying 10% Tax Rate to Assessed Tax

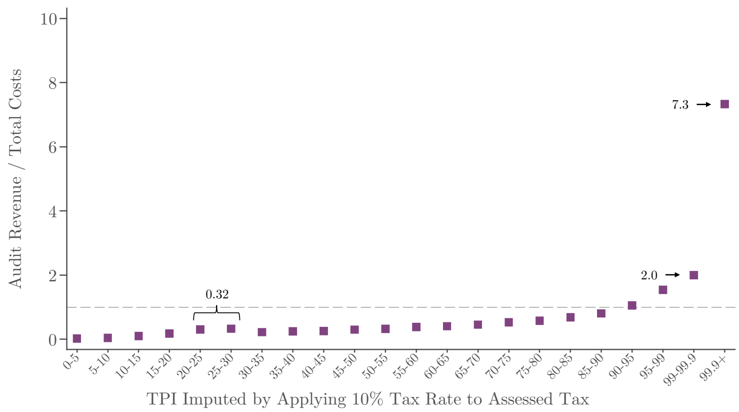

B. TPI Imputed by Applying 20% Tax Rate to Assessed Tax

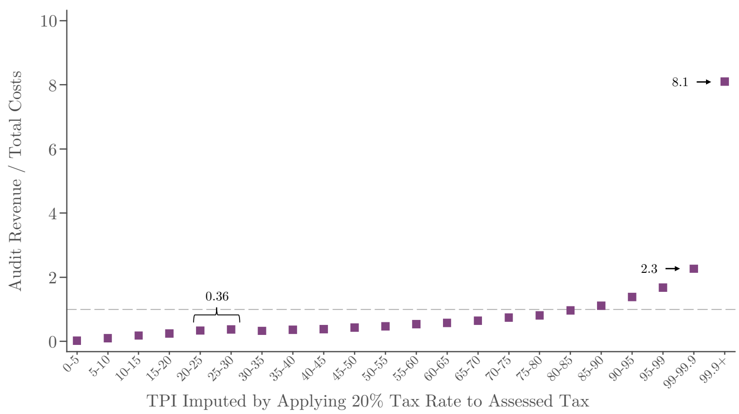

Notes: This figure repeats Figure II (Panel B) using imputed post-audit TPI on the horizontal axis instead of pre-audit reported TPI. We impute post-audit TPI by dividing the change in assessed tax by a 10% average marginal tax rate (Panel A) and by a 20% average marginal tax rate (Panel B) and adding the change to reported TPI.

APPENDIX FIGURE A.VIII:

Average Hours per In-Person Audit and Auditor Wage Rate, by Income Group

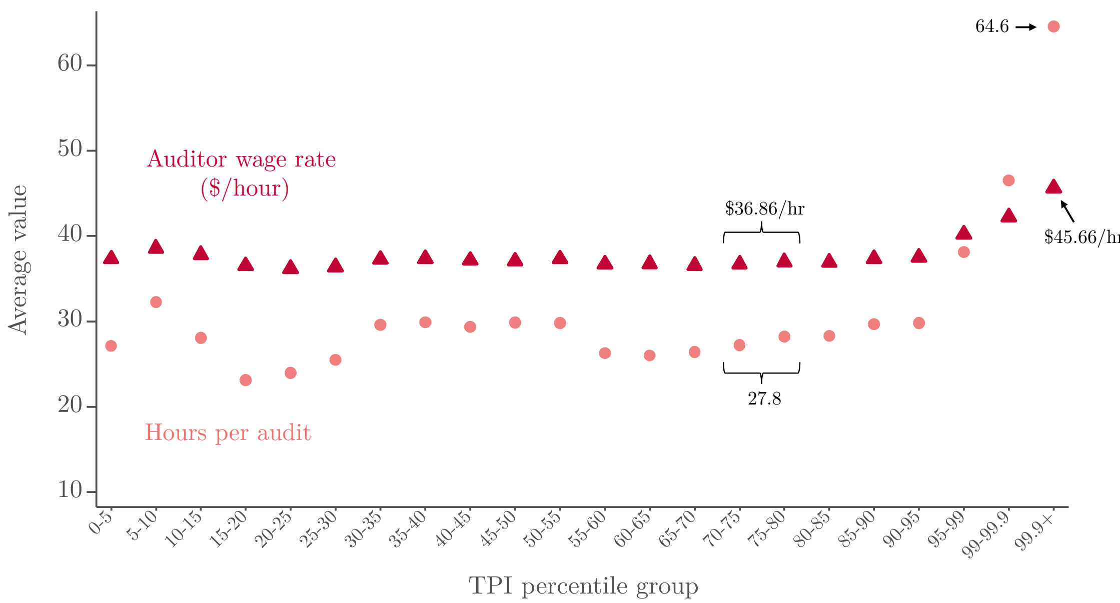

Notes: This figure presents the average values of the components of labor costs (auditors' wages times hours spent on exam) for in-person audits of tax returns filed for tax years 2010–2014 by the taxpayer's total positive income (TPI). The x-axis groups TPI into bins of five percentiles and splits out the top bin into the 95–99th and 99–99.9th percentiles and the top 0.1%. Average hours per audit include projected labor hours accrued after the observed 7–11 year post-audit sample window.

APPENDIX FIGURE A.IX:

Marginal Revenue

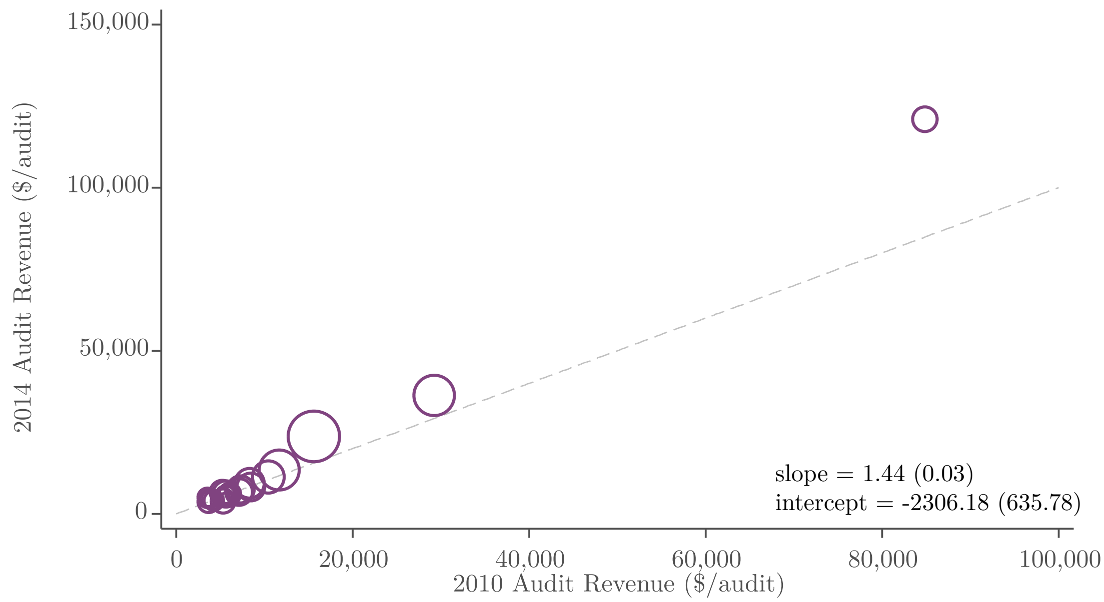

Notes: This figure shows average revenue raised per in-person audit for each 5-percentile TPI bin for audits of returns from tax year 2014 = against the values for the same TPI bins for audits of returns from tax year 2010. Average revenues include projected revenue collected and costs accrued after the observed 7 and 11 year post-audit sample window for 2014 and 2010 audits respectively. Revenues are discounted using a 3% discount rate since average revenues lag average costs by about one year on average. In particular, we use data from the 2003 tax year to separately discount the revenues raised and costs accrued each year post-audit back to the tax year. We then use the ratio of the discounted series (net present value of revenues divided by costs) to adjust revenues downwards to align the two paths.

APPENDIX FIGURE A.X:

Audit Probability, Revenue Collected and Wage Costs per NRP Audit, by Year

A. Audit Probability, Revenue Collected and Wage Costs

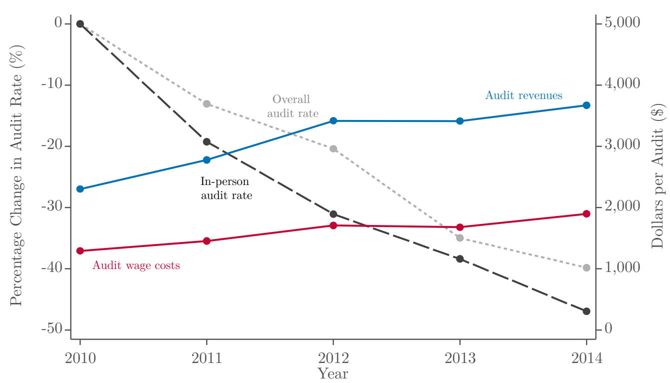

B. Components of Wage Costs

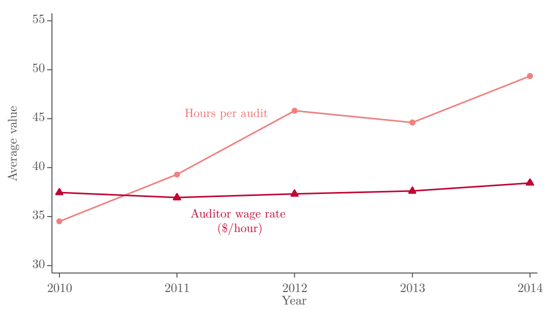

C. Audit revenue over Wage Costs

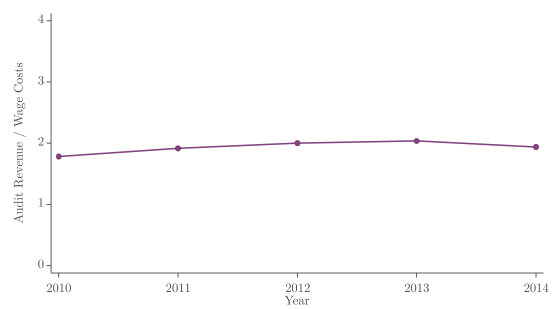

Notes: Panel A presents the percentage change in overall and in-person audit rates, total revenues raised, and direct labor costs (auditors' wages times hours spent on exam) per National Research Program (NRP) random audit for NRP study tax years in our sample frame (2010–2014). Panel B shows each component of labor costs (auditors'wages and hours worked per audit) by year. Panel C shows average revenue per audit divided by costs per audit for each NRP study tax year. Total revenue is the sum of additional tax liability, penalties and interest collected. Average costs and revenues include projected costs incurred and revenue collected after the observed 7–11 year post-audit sample window. Revenues are discounted using a 3% discount rate because revenues lag costs by about a year on average. In particular, we use data from the 2003 tax year to separately discount the revenues raised and costs accrued each year post-audit back to the tax year. We then use the ratio of the discounted series (net present value of revenues divided by costs) to adjust revenues downwards to align the two paths.

APPENDIX FIGURE A.XI:

Average Revenue and Costs per NRP Audit, by Income and Year

A. Average Audit Revenue

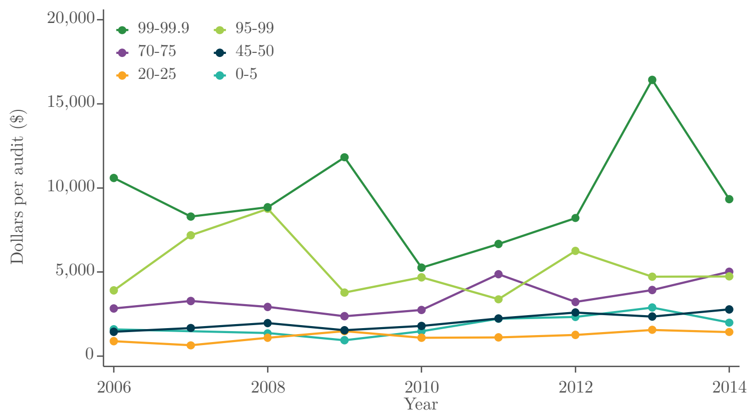

B. Average Costs

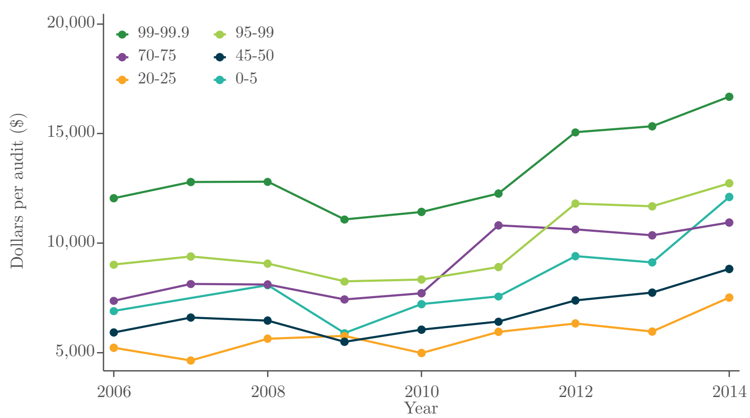

Notes: This figure presents total audit revenues raised (Panel A) and total costs accrued (Panel B) per random audit in each National Research Program (NRP) study tax year for select total positive income (TPI) percentile bins. Total revenue is the sum of additional tax liability, penalties and interest collected. Total costs are the sum of labor costs (auditors' wages times hours spent on exam) and additional costs (labor/fringe/primary, organization-wide, and overhead/HQ costs), which are allocated in proportion to direct labor costs. Average costs and revenues include projected costs incurred and revenue collected after the observed 7–11 year post-audit sample window. Revenues are discounted using a 3% discount rate because revenues lag costs by about a year on average. In particular, we use data from the 2003 tax year to separately discount the revenues raised and costs accrued each year post-audit back to the tax year. We then use the ratio of the discounted series (net present value of revenues divided by costs) to adjust revenues downwards to align the two paths.

APPENDIX FIGURE A.XII:

Within-Taxpayer Impact of Audits on Future Tax Payments, by Presence of Income

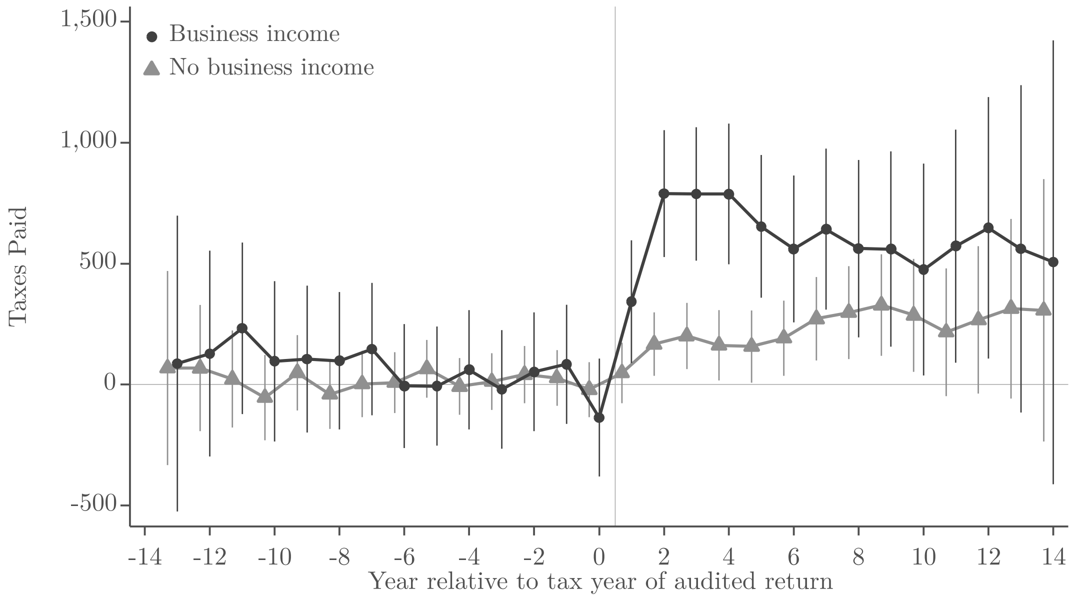

Notes: This figure presents estimates of the change in taxes paid each year post-audit for individuals selected for random audit by the National Research Program (NRP) separately for individuals with and without business income (as measured by income on Schedule C, E and F). The control group is a matched sample of individuals not selected for random audit. Collected tax revenue is winsorized at the 99th percentile of the population distribution to limit the influence of outliers. The plotted estimates show the difference in taxes paid between control and treated individuals in each year in a single difference specification.

APPENDIX FIGURE A.XIII:

Within-Taxpayer Impact of Audits on Future Tax Payments, by NRP Study Year

A. 4 year groupings

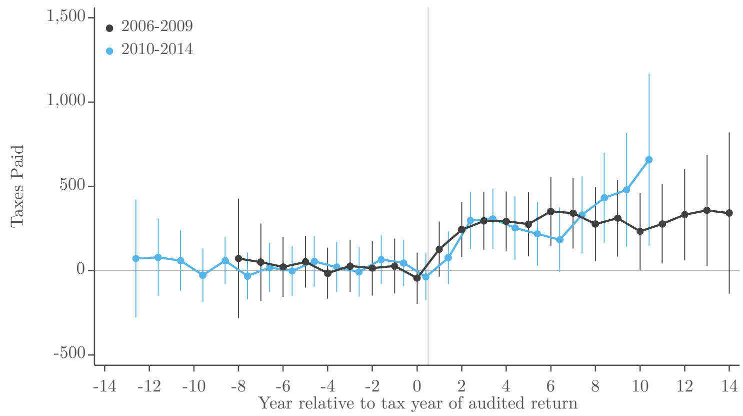

B. 2–3 year groupings

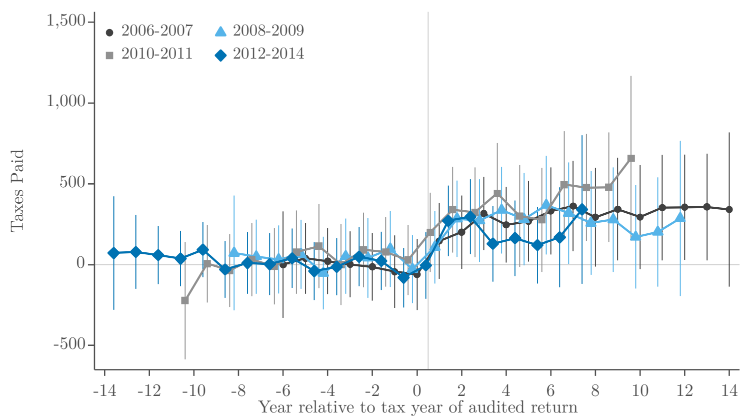

Notes: This figure presents estimates of the change in taxes paid for each year post-audit for the full population of individuals selected for random audit by the National Research Program (NRP), by groups of NRP study tax years. The control group is a matched sample of individuals not selected for random audit. Collected tax revenue is winsorized at the 99th percentile of the population distribution to limit the influence of outliers. The plotted estimates show the differences in taxes paid between control and treated individuals in each year in a single difference specification.

APPENDIX FIGURE A.XIV:

Within-Taxpayer Impact of (Non-Random) In-Person Audits on Future Tax Payments

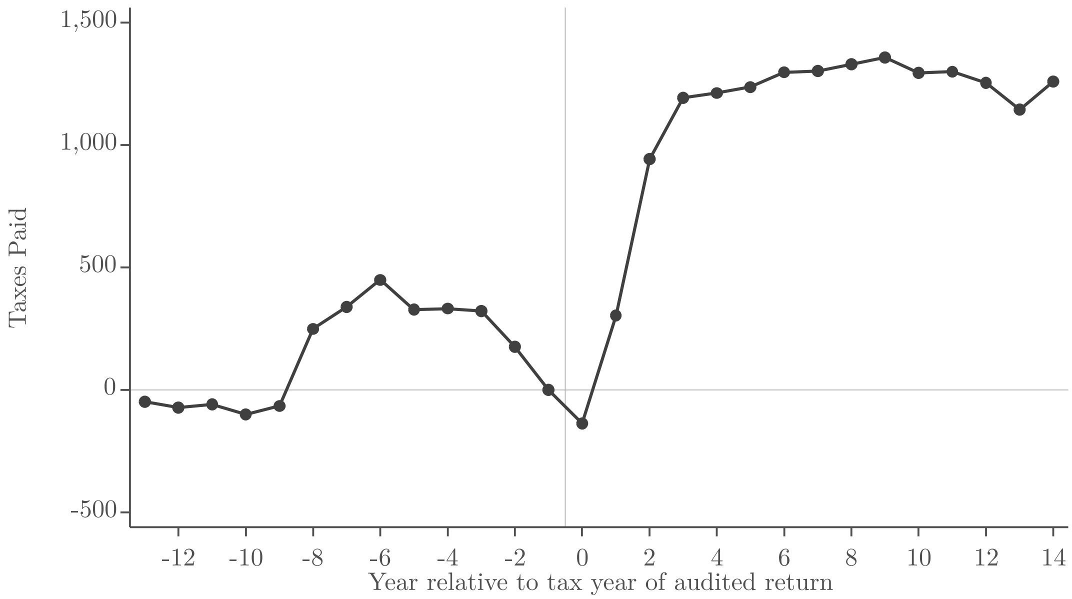

Notes: This figure presents estimates of the change in taxes paid in each year post-audit for individuals selected for an in-person audit. The control group is a matched sample of individuals not selected for audit. Coarsened matching is done based on income, lagged income, and return characteristics. The plotted estimates show the result of a difference-in-differences comparison in taxes paid. The figure compares treated and control individuals, comparing both to their respective taxes paid in the year before the audit. Collected tax revenue is winsorized at the 99th percentile of the population distribution to limit the influence of outliers.

APPENDIX FIGURE A.XV:

Deterrence Effect over Initial Audit Revenue (Winsorized and Unwinsorized), by Income

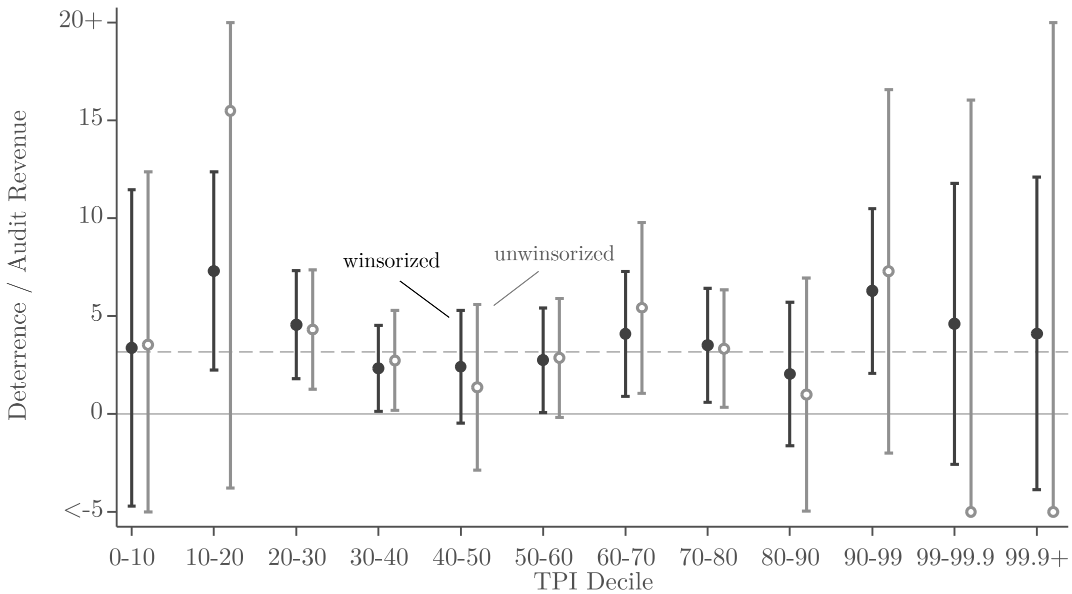

Notes: Panel A presents estimated deterrence effects by the taxpayer's Total Positive Income (TPI) decile. Deterrence effects are the net present value (NPV) of total additional taxes paid post-audit divided by the NPV of upfront revenue raised per National Research Program (NRP) random audit. Additional taxes paid are estimated using a matched differences-in-differences specification, which compares taxes paid for each year post-audit by individuals selected for random audit relative to a matched sample of individuals not selected for random audit. This specification is run separately by TPI decile. Collected tax revenue is winsorized at the 99th percentile of the population distribution to limit the influence of outliers.

APPENDIX FIGURE A.XVI:

Taxpayer Burden of Audits, by Income

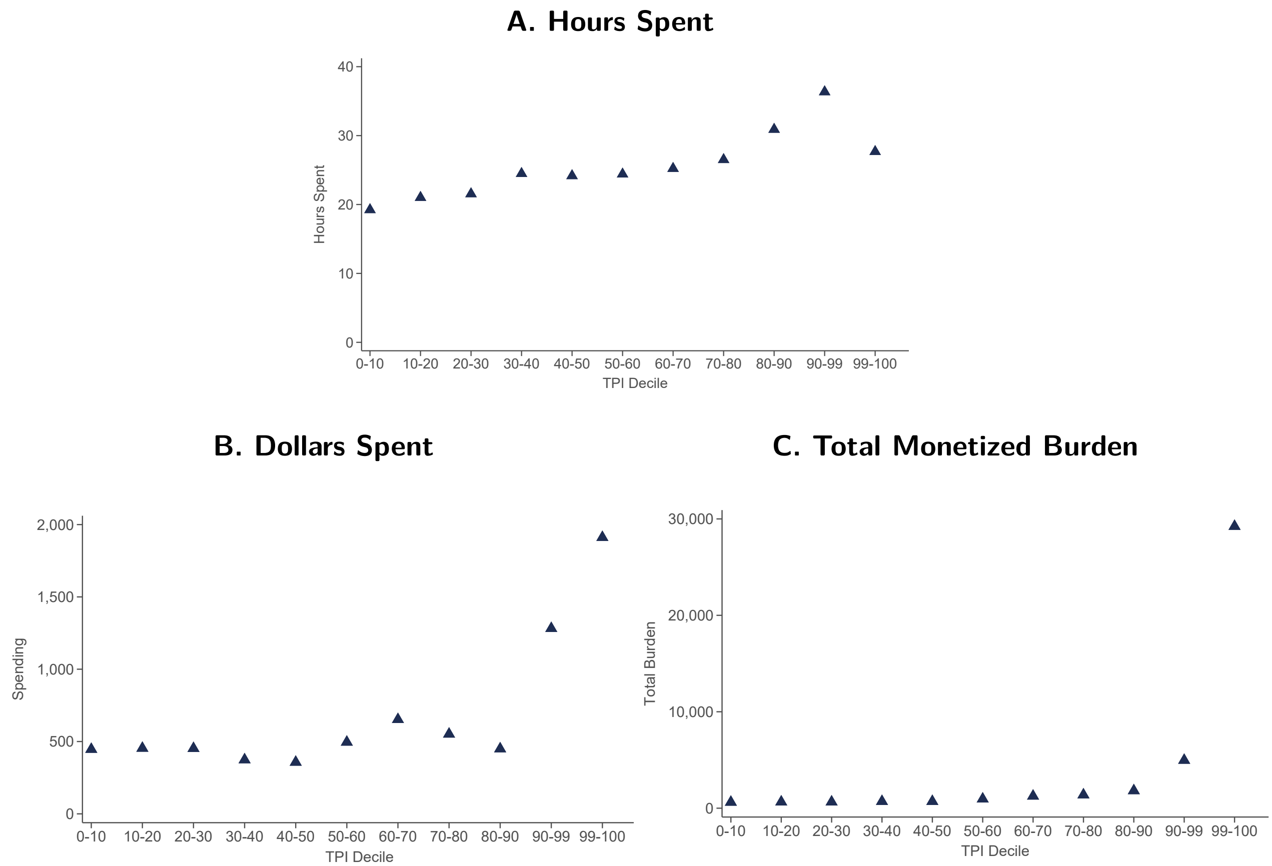

Notes: This graph presents the taxpayer burden of audits using a 2023 representative survey of audited taxpayers conducted by the IRS. Taxpayers are asked about the time and money they spent to comply with the audit. Responses are matched to taxpayer TPI using coarse bins corresponding closely to TPI percentile thresholds. The figures report the average hours spent (Panel A) and dollars spent (Panel B) for each audit. In panel C, the total monetized burden imputes an hourly wage for the taxpayer by dividing their total income by roughly 2000.

APPENDIX FIGURE A.XVII:

Fraction of Audits with No Additional Assessed Tax Liability, by Income

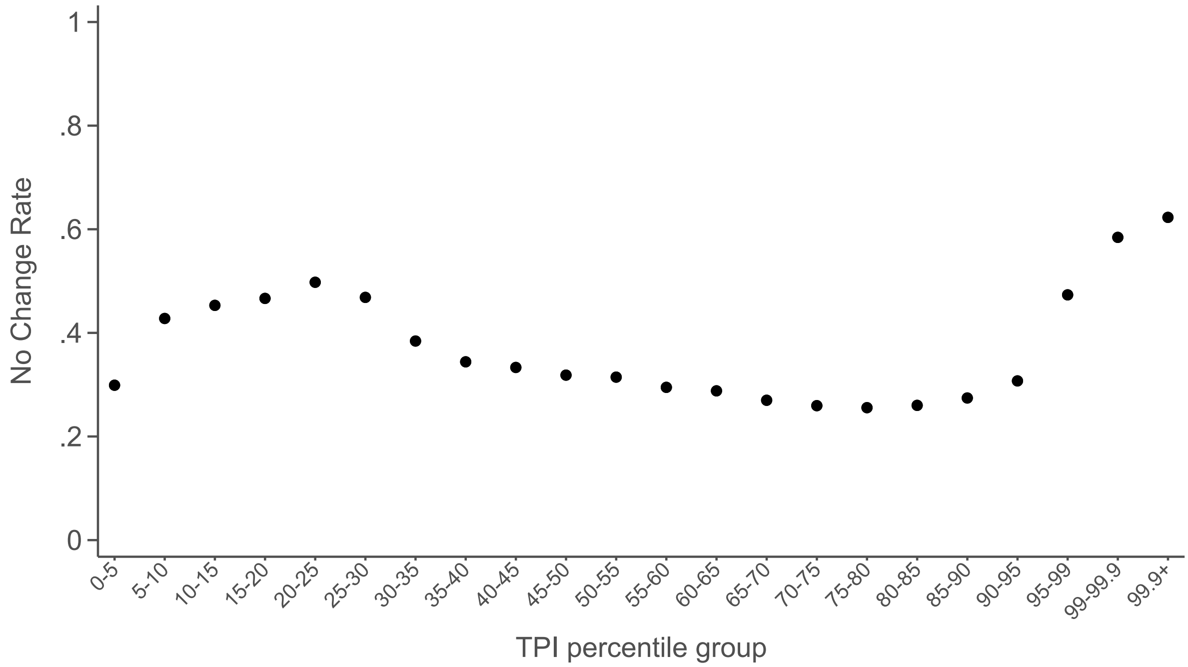

Notes: This graph presents the fraction of in-person audits with no additional assessed tax liability, separately by TPI bin.

APPENDIX FIGURE A.XVIII:

Taxpayer Burden of Audits, by Income and Whether Additional Tax was Assessed

A. Hours Spent

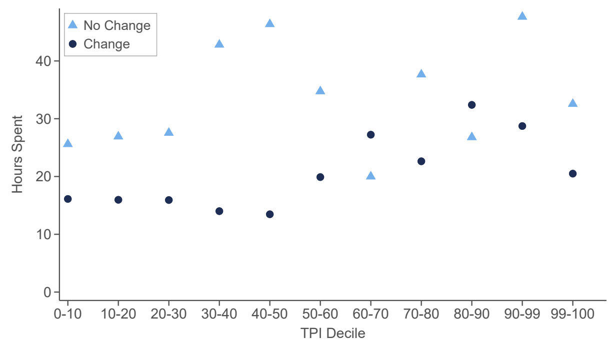

B. Dollars Spent

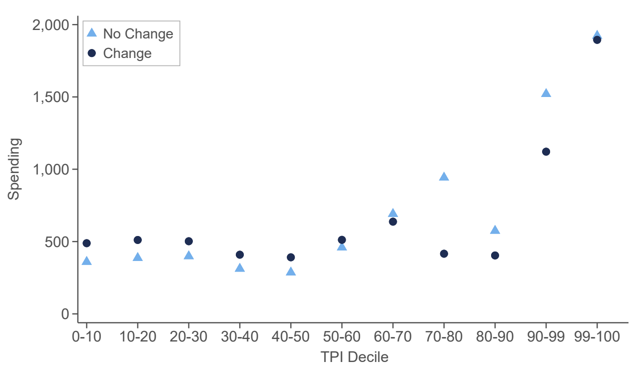

C. Total Burden

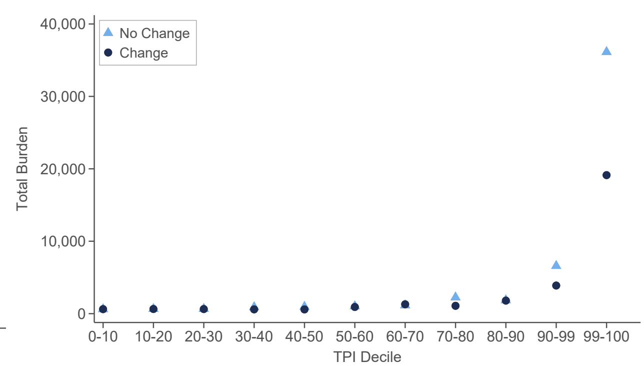

Notes: This figure repeats Appendix Figure A.XVI but reports the estimates separately by whether the audit resulted in additional assessed tax liability.

APPENDIX FIGURE A.XIX:

Response of EITC Amount Claimed to Random Audits, by Whether the Auditor Recommended Reducing the EITC Amount

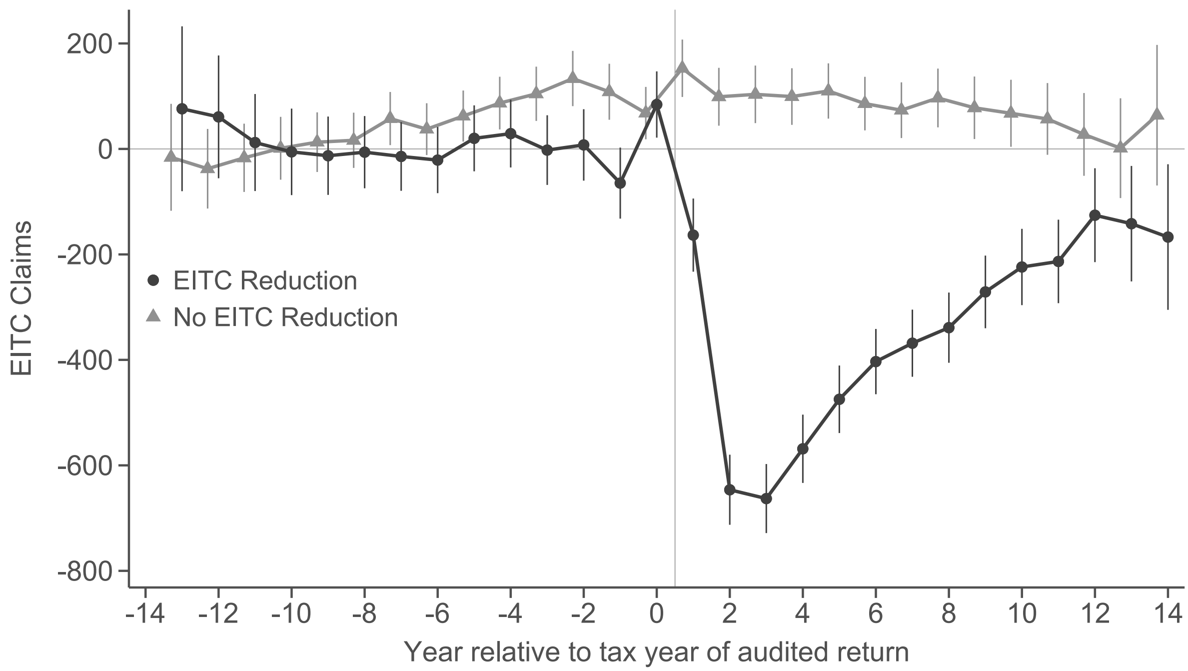

Notes: This figure presents differences in average EITC amounts claimed between taxpayers subject to random audit by the NRP and matched controls. It restricts to randomly audited taxpayers who claim the EITC in the tax year selected for random audit and their matched controls, who also claim the EITC and have similar incomes in that tax year. The sample is further split by whether the randomly audited taxpayer's auditor recommended reducing the EITC amount claimed on the return. Each subgroup is then compared to their income-and-EITC-claim-status-matched controls. The difference for those with an auditor-recommended reduction in EITC amount is shown with circles and for those without an auditor-recommended reduction in EITC amount in triangles.

B Data Appendix

This appendix discusses the data used in our analysis and the methods used to compute the costs and revenues associated with each audit. As noted in the main text, we use two types of internal IRS data as well as data from a survey on taxpayer burdens. In what follows, we detail how we handle each type of data to estimate revenues and costs.

B.1 Audit-Level Enforcement Data: Revenue and Direct Labor Cost Estimates

We use administrative, audit-level data from the IRS's internal enforcement database. The transaction-level dataset records all activities associated with a given audit (e.g. time spent by the auditor, payments collected from the taxpayer, etc). These data allow us to calculate revenues collected for each audit and direct labor costs expended on each audit.

We begin with all rows in the enforcement database where the taxpayer is an individual. [1] This includes both operational exams (in-person and correspondence) as well as random NRP audits.

Identifying a single audit. We combine examinations conducted at the same time of multiple tax returns filed by the same taxpayer into a single audit to capture scenarios where an auditor makes similar adjustments to returns for multiple years. We join examinations that start between a prior examination's start date and 90 days after the prior examination's end date. For example, if a taxpayer's 2016 return is examined from January 2018 to January 2019 and an examination begins of the same taxpayer's 2015 or 2017 tax return between January 2018 and April 2019, we sum the revenue and costs associated with that return as part of the revenues and costs of the examination of the 2016 return. By this definition, 81.9% audits are associated with one tax return, 12.5% are associated with 2 tax returns, and 5.6% are associated with 3 or more tax returns.

Date variables. We estimate the date an audit started using the transactions associated with the audit. In addition, we use five variables included in the enforcement data to estimate four year variables that help us track each audit over time, adjust for inflation and the timing of costs and revenues, etc.. We define the audit start year as the year in which transactions associated with an audit began. We define the primary year to be the tax year of the return that triggered the enforcement process by our definition above.[2] For example, suppose that an audit of a tax year 2014 return (which would have been filed in calendar year 2015) began in calendar year 2017. The audit start year would be 2017 and the primary year would be 2014. The labor year is the fiscal year for which the hours are recorded.[3] We use the labor year to determine the auditor's wage rate. Finally, the revenue year is the fiscal year in which payments are recorded in the transaction data.

Identifying the stage of the audit. We assign revenue and costs by the stage of the audit (i.e. exam, appeals, collections). To do so, we use the function code associated with each transaction to classify each transaction into these stages.

Assigning TPI percentiles. We assign individuals to the income distribution using the taxpayer's total positive income (TPI) in the tax year that triggered the audit (i.e., the "Primary Year" and

[1] More than 99% of these are audits of individual income tax returns, with the remainder including income tax returns of trusts and gift or estate tax returns.

[2] For any audit identified as a random NRP audit, we use the NRP study year for the "primary year". This changes the value of primary year in 8,000 transaction-level rows (out of 295,000 identified by the NRP data).

[3] When available, labor year is set equal to the fiscal year of the assessment. If assessment fiscal year is missing, we use the year after the exam start year. When both assessment fiscal year and exam start year are populated, assessment fiscal year is one greater than exam start year in 68.8% of rows (and equal in 22.6% of rows). Exam start year may differ from assessment fiscal year in situations where there was a delay between the beginning of the examination and when the work took place. This leaves very few cases with a missing value for labor year. In those few cases, there are no associated hours.

by-year, population-level percentiles of reported TPI. For example, if the primary year of an audit was 2012, we determine that audit's TPI percentile using population-level TPI percentiles from 2012.

Estimating revenue. To estimate the revenue from a given audit, we add total enforcement revenue from taxes, penalties, and interest for the following stages:

1. Exam
2. Appeals and Counsel
3. Collections: notices (1st, 2nd, 3rd, and 4th notices)
4. Collections: ACS
5. Collections: any revenue collected while in the queue for field collections
6. Collections: field collections

Total revenue from collections is the sum of items 3 through 6 above, and total revenue from an audit is the sum of revenue from the exam, appeals, and collections stages. Before summing, all revenue variables are adjusted for inflation using the CPI-U-RS to 2016 dollars.

Estimating direct labor costs (exam and appeals stages). Each row of data notes the hours spent by the IRS employee as well as their General Scale (GS) pay grade. We estimate direct labor costs for the exam and appeals stages by multiplying the hours in a given transaction by the matched hourly GS pay rate (described below). We note that hours and costs associated with legal counsel are listed as a separate category of transactions. Because these expenses are generally incurred during the appeals phase, we add the direct costs from any "counsel" hours to appeals to get a total direct labor cost estimate for appeals and counsel. We then use the year in which the activities and expenses were incurred (i.e. "labor year" defined above) to adjust these costs for inflation.

Estimating direct labor costs (field collections). While hours and GS grade are stored for cases that are sent directly to collections, the hours and GS grade information for collections personnel is not stored in the enforcement database for cases that originated in the exam stage.

To estimate the direct labor cost of field collections, we use transactions related to cases that went straight to collections (and therefore have associated hours and GS grade information) from the enforcement database. We estimate direct labor costs for these cases in the same way as exam and appeals cases in our main dataset. We create a 10-by-10 index, where one axis plots deciles of total positive income (TPI; the lowest decile restricted to zero) and the other axis plotting deciles of amount assessed. Each cell contains the average cost estimate associated with that combination of TPI-decile and amount assessed-decile. For each audit in our primary dataset, we determine the relevant decile along both dimensions and apply the corresponding cost estimate.

General Schedule (GS) pay rates. We determine the relevant hourly GS pay rate using the year and location in which the labor activities took place. We proceed as follows. First, we determine the relevant zip code for each individual in our dataset. We determine zip code for a given tax year first by using the modal zip code for an individual's third-party information reporting in that tax year. If there is no zip code found that way, we use the zip code from the individual's 1040. If there is still no zip code, we use the modal third-party information reported zip code from the previous tax year, then the previous tax year's 1040, going back 5 tax years. If there is still no zip code, we apply the average location adjustment from the matched zip codes (matched as described below).

We create a mapping of zip code to FIPS codes, and then FIPS codes to the localities provided in the historical GS pay scale data. After creating this mapping, we merge location-specific hourly pay rates to the transactions the enforcement data by year, location, and GS grade. We use the GS pay scale hourly rate for Step 5. If we had a matched zip code in the enforcement data but no associated GS pay rate for that zip code, we applied the Rest of US (RUS) rate for that GS grade and year. [4]

We then take an average of the difference between the assigned payrates and the base GS payrates to determine our average location adjustment.[5] We use this average location adjustment for cases where there was no matched zip code.

B.2 Accounting Data: Non-Direct Labor Cost Estimates

Our second data source is internal line-item accounting data from the IRS. These data enable us to include all potential costs associated with audits beyond the direct labor hours spent by the auditors.

As background, the IRS has four large operating divisions that deal with taxpayers: (1) Small Business/Self-Employed (SB/SE), (2) Large Business and International (LB&I), (3) Wage and Investment (W&I), and (4) Tax Exempt and Goverment Entities (TE/GE). These operating divisions are responsible for different populations of taxpayers (Internal Revenue Manual 1.1.1).

SB/SE is the operating division responsible for audits of individual tax returns and therefore the operating division relevant for this project.[6] There are three organizations within SB/SE: collection, examination, and operations support (Internal Revenue Manual 1.4.40.2). The examination organization is responsible for both field examination and correspondence examination (Internal Revenue Manual 1.1.16.5).

The internal accounting data we use is organized by enforcement function (e.g., field exam versus correspondence exam) rather than by operating division (e.g., SB/SE versus LB&I). We use these data to construct cost measures for five different enforcement functions: Field Exam (i.e. in-person exams), Correspondence Exam, Field Collections, Collections Notices, and the Automated Collection System (ACS). We then use these numbers to calculate the average cost of each audit by assuming that total costs associated with the audit are proportional to the labor costs associated with the audit. Where possible, we construct our multipliers using SB/SE-specific cost information.

In-person exams. To calculate the cost multipliers for in-person exams, we use internal IRS accounting data for the Field Exam organization from 2011–2020. These data include line-item level cost information for all Field Exam operations as well as total costs for Field Exams conducted by SB/SE. We create our in-person exam multipliers by calculating the appropriate ratios of these line-items to total costs using the information available for all Field Exam costs and applying those ratios to the total Field Exam costs for SB/SE. We describe this process in detail below.

First, we determine the total costs to the government associated with in-person exams. This is similar to the total costs given in the internal accounting data with two slight differences. First, the total costs given in the Field Exam accounting data exclude "imputed" costs for support services provided by other government agencies and not directly part of the IRS's budget. For example,

[4] The historical GS pay scale data is available at `https://www.opm.gov/policy-data-oversight/pay-leave/salaries-wages/`.

[5] This average is weighted by number of hours.

[6] SB/SE audits both individuals and small businesses such as partnerships, S-Corporations, and C-Corporations with assets under $10 million. The same revenue agents conduct audits of complex individual returns and business returns. C-Corporations with assets greater than or equal to $10 million are handled by LB&I. W&I conducts pre-refund examinations of EITC returns, which are not included in our data. For more information on the SB/SE operating division, visit `https://www.irs.gov/about-irs/small-business-self-employed-division-at-a-glance` .

Appendix Table A.I:
Total Costs for Field Exam ($million)

Year	Total Field Exam Costs	Imputed %	Appeals %	Total Gov. Costs	SBSE Total Gov. Costs	SBSE Direct Labor Costs
2011	3,224	8.34%	3.44%	3,382	1,722	314
2012	3,201	8.35%	3.45%	3,358	1,645	325
2013	3,158	9.39%	3.33%	3,349	1,689	319
2014	3,091	8.97%	3.22%	3,268	1,652	284
2015	2,955	8.14%	0.29%	3,188	1,547	294
2016	2,787	9.78%	0.40%	3,048	1,635	265
2017	2,697	8.86%	0.27%	2,928	1,549	256
2018	2,650	11.91%	0.25%	2,959	1,558	234
2019	2,568	11.85%	0.23%	2,867	1,523	200
2020	2,475	10.74%	0.19%	2,736	1,363	132

the Department of Agriculture runs payroll for several government agencies, and the Bureau of the Fiscal Services processes payments for the IRS. While these costs are not relevant for internal IRS budgeting purposes, they contribute to the total fiscal cost of audits. We observe imputed costs and include them in the total costs of in-person exams. Second, the internal accounting data include the cost of appeals in their total costs for in-person exams. We separate exams and appeals in the labor input process and construct separate measures of the direct labor costs of appeals. We exclude these costs from our calculations for in-person exam and use them to construct a separate cost multiplier for the appeals process. We next determine the total costs associated with audits for the SB/SE operating division. We do not observe imputed costs and appeals by operating division, so we assume that the fraction of the budget that consists of imputed costs is the same for each operating division. In practice, this means we can apply the ratios from the overall costs for Field Exams to the total costs from SB/SE. Table A.I shows the total costs to the government for SB/SE field exams.

After determining the total costs associated with audits for the SB/SE operating division, we split these costs into three broad buckets: (1) non-direct labor-related costs, (2) organization-wide costs, and (3) other overhead. We use the line-item level data for all field exams to assign costs to categories. Non-direct labor-related costs include time spent by auditors on tasks other than examining returns (e.g. training), fringe benefits, management labor costs, training costs, and other primary labor-related costs. Organization-wide costs include space/rent and information technology costs incurred by auditors. Finally, overhead costs are allocations of accounting costs from the central IRS management that overseas the audit programs.[7] We only observe these line-item costs totaled across all of field exam, including LB&I and TE/GE field exam in addition to SB/SE field exam. We again assume that these line-items are a constant fraction of the total budget for each operating division within field exam. In Table A.III, we calculate what percentage of total relevant costs comes from each of these three categories and apply these percentages to the total relevant costs estimated for SB/SE to estimate the value of each of these categories of costs for SB/SE. This gives us the costs for SB/SE that will be used in the numerator of our cost multipliers

[7] Table A.II lists the specific line-items included in each of these broad categories.

Appendix Table A.II:
Categories for Field Exam Costs

Non-Direct Labor-Related Costs	Organization-Wide Costs	Other Overhead
Labor	Rent /Building	Internal BU
Benefits	IT	Imputed Cost
Services/Supplies	Printing/Postage	
Traveling	TAS	
Training	TE/GE	
Enforcement	Depreciation	
Printing	Appeals	
Moving Expense	Other Finance	
ADP Operations	Workers Comp	
Space & Housing	UCFE	
Rent	HCO	
Equip.-Non-ADP	WISK Other	
Postage	Corporate S&F	
Communications	Comm & Liaison	
Misc. Revenue	SB/SE	
Misc. Expense	LB&I	
	Corr. Exam Support	

for field exam.

The last piece of information we need to calculate our multipliers is an estimate of total direct labor costs from SB/SE Field Exam. This enables us to calculate the ratio of the different types of costs discussed in Tables A.II and A.III to the direct labor costs we estimate using the administrative enforcement data.

To calculate total direct labor costs for SB/SE, we proceed analogously to our measure of direct labor costs for individual audits measured above, but we now pull all transactions from the administrative enforcement data associated with in-person exams and with the SB/SE operating division (i.e. including businesses). We calculate total direct labor costs as described above for our audit-level dataset. For transactions related to businesses instead of individuals, we apply an average location adjustment rather than a zip-code specific hourly pay rate.[8] The total direct labor costs from SB/SE in-person exams is given in the last column of Table A.I.

We calculate three multipliers for in-person exams using the three broad cost categories described above: (1) non-direct labor-related costs per-dollar of direct labor costs,[9] (2) organization-wide costs

[8]To check our direct labor cost estimates for all of SB/SE field exam, we compared the underlying counts of hours with headcount estimates of the total number of active in-the-field Revenue Agents and Tax Compliance Officers for 2016–2018. The estimates imply that active field examiners average about 70% of a 40-hour work week directly on exams. This value is consistent with conversations we had with individuals who work in the SB/SE field exam unit. We do not include 2019 and 2020 in this exercise. There was an influx of new examiners in 2019. As a result, averages are likely depressed by training time. The pandemic led IRS to pause operations in 2020; any hours estimates would not be reflective of usual IRS operations.

[9]When calculating the multiplier for non-direct labor-related costs, we subtract our direct labor costs estimate

Appendix Table A.III:
Cost Breakdown for Field Exam

Year	Non-Direct Labor Costs	Organization-Wide Costs	General Overhead Costs
2011	58.39%	21.80%	19.80%
2012	56.63%	22.14%	21.22%
2013	52.74%	18.24%	29.02%
2014	52.66%	19.00%	28.34%
2015	51.52%	13.10%	35.39%
2016	51.39%	12.64%	35.96%
2017	52.90%	11.47%	35.63%
2018	50.32%	11.20%	38.48%
2019	50.78%	11.63%	37.59%
2020	53.70%	12.01%	34.29%

per dollar of direct labor costs, and (3) other overhead costs per dollar of direct labor costs (see Table A.IV). We sum these three values for an overall multiplier.

Because we focus on audits of returns for tax years 2010–2014 (which are filed in early 2011–2015), we use the average of the 2011–2015 fiscal year values of these multipliers. The overall value is 4.39. That is, for every dollar of direct labor costs spent on in-person exams, we include $4.39 of non-direct labor costs, organization-wide costs, and other overhead costs. Non-direct labor costs account for nearly half of the additional costs.

Correspondence exams. To calculate the costs for correspondence exams, we use internal IRS accounting data for the Correspondence Exam organization from 2011–2020. We obtain data on costs for all correspondence exams within SB/SE, not just those of individuals. We create our correspondence exam multipliers by calculating the appropriate ratios with the detailed data for all correspondence exams and applying those ratios to the total costs for SB/SE. We describe this process in detail below.

First, we determine the total costs of conducting correspondence audits by adding in the imputed costs that are incurred by other government agencies on behalf of Correspondence Exam activities. We next determine the total relevant costs for correspondence exam for the SB/SE operating division. Because imputed costs by operating division are not included in the data, we assume that the fraction of costs that are imputed costs is the same across operating divisions within correspondence audits. This means that we can apply the ratios from the overall costs for Correspondence Exams to the total costs from SB/SE. This is shown in Table A.V.

Table A.VIII also shows how we apply the relevant ratios to the total costs for SB/SE to estimate the division of costs between our three categories.[10]

form the numerator.

[10]Before we break these relevant total costs down into our three cost subcategories, we need to subtract the component of the Correspondence Exam efforts that were conducted in support of field exam audits. These costs were included in the costs of Field Exam, because they are costs incurred as a result of the Field Exam program, and therefore should be excluded from the total costs for Correspondence Exam. We assume these costs in support of field exam efforts are incurred proportionally across the line items and subtract them from our cost categories accordingly.

Appendix Table A.IV:
Cost Multipliers per Direct Labor Dollar for In-Person Exam

Year	Non-Direct Labor Cost Multiplier	Organization-Wide Cost Multiplier	General Overhead Cost Multiplier	Overall Cost Multiplier
2011	2.21	1.20	1.09	4.49
2012	1.87	1.12	1.07	4.06
2013	1.79	0.96	1.54	4.29
2014	2.06	1.10	1.65	4.81
2015	1.71	0.69	1.86	4.27
2016	2.17	0.78	2.22	5.16
2017	2.20	0.69	2.15	5.05
2018	2.35	0.74	2.56	5.65
2019	2.88	0.89	2.87	6.64
2020	4.54	1.24	3.54	9.32
2011-2015 Average	**1.93**	**1.02**	**1.44**	**4.39**

Appendix Table A.V:
Total Costs for Correspondence Exam ($million)

Year	Total Corr. Exam Costs	Imputed %	Total Gov. Costs	SBSE Total Gov. Costs	SBSE Direct Labor Costs
2011	430.0	12.88%	485.4	248.1	16.7
2012	476.9	12.42%	536.1	252.7	16.5
2013	444.8	15.90%	515.5	249.5	16.0
2014	463.0	14.68%	531.0	257.3	12.6
2015	392.7	15.18%	452.3	224.2	11.8
2016	377.8	16.16%	438.9	223.3	11.5
2017	357.8	14.87%	411.0	171.8	10.0
2018	350.8	19.86%	420.5	189.7	9.7
2019	347.3	19.83%	416.2	179.1	8.6
2020	362.3	18.40%	428.9	182.9	5.7

Appendix Table A.VI:
Cost Breakdown for Correspondence Exam

Year	Non-Direct Labor Costs	Organization-Wide Costs	General Overhead Costs
2011	49.34%	32.35%	18.32%
2012	44.59%	29.96%	25.45%
2013	43.75%	12.56%	43.69%
2014	41.01%	11.64%	47.36%
2015	44.37%	17.84%	37.78%
2016	42.36%	15.32%	42.32%
2017	44.35%	17.10%	38.54%
2018	41.47%	15.47%	43.05%
2019	42.45%	16.49%	41.06%
2020	42.55%	16.68%	40.76%

We calculate three multipliers for correspondence exams using the three broad cost categories described above: (1) non-direct labor-related costs per-dollar of direct labor costs,[11] (2) organization-wide costs per dollar of direct labor costs, and (3) other overhead costs per dollar of direct labor costs (see Table A.VII). We sum these three values for an overall multiplier.

We use the average of the 2011–2015 values of these multipliers. The overall value is 16.07. That is, for every dollar of direct labor costs spent on in-person exams, we include $16.07 of non-direct labor costs, organization-wide costs, and other overhead costs. For correspondence exams, non-direct labor costs account for about 40% of these additional costs.

Appeals. Both field and correspondence exams recommendations can be appealed, which sends the case to the IRS's Independent Office of Appeals. To calculate the cost of the appeals stage of audits, we take the line-item costs for Appeals from the Field Exam data and divide by direct labor costs for appeals and counsel from the administrative enforcement data.[12] Neither the internal accounting data nor the administrative enforcement data splits appeals costs by operating division. Consequently, we use the internal accounting data for all field exams and all transactions associated with appeals and counsel from the administrative enforcement data to estimate the direct labor costs. We use the average from 2011–2015, which is 0.57.[13]

[11] When calculating the multiplier for non-direct labor-related costs, we subtract our direct labor costs estimate form the numerator.

[12] There is not a corresponding Appeals line-item from the Correspondence exam data, perhaps because appeals of Correspondence exams are even rarer than appeals of in-person exams.

[13] In 2015, there was a marked drop in the reported line item for appeals in the internal IRS accounting data. The drop is extensive enough that, for 2015, our estimated direct labor costs for appeals from the administrative enforcement data (which do not experience the same drop) are an order of magnitude bigger than the appeals line item listed in the internal IRS accounting data. We presume this drop in costs in the accounting data is the result of an internal change in how certain costs were allocated across the field exam operating division. As a result, the "missing" costs for appeals in 2015 would appear in a different line item and would be included in our overall cost calculations through that channel. This means that we are slightly underestimating the average overhead costs of

Appendix Table A.VII:
Cost Multipliers per Direct Labor Dollar for Correspondence Exam

Year	Non-Direct Labor Cost Multiplier	Organization-Wide Cost Multiplier	General Overhead Cost Multiplier	Overall Cost Multiplier
2011	6.34	4.81	2.73	13.88
2012	5.85	4.60	3.91	14.36
2013	5.82	1.96	6.81	14.60
2014	7.40	2.38	9.70	19.49
2015	7.44	3.39	7.18	18.01
2016	7.23	2.98	8.22	18.43
2017	6.62	2.94	6.62	16.17
2018	7.08	3.02	8.39	18.49
2019	7.82	3.43	8.53	19.78
2020	12.61	5.34	13.04	30.98
2011-2015 Average	**6.57**	**3.43**	**6.07**	**16.07**

Collections. Not everyone pays what they owe, even after the amount owed is not in dispute. The IRS collections process begins with notification letters to the taxpayer indicating that they have an unpaid balance due. If the taxpayer does not respond to the notifications, the case will be handled by the Automated Collection System (ACS) or by a local field office (Field Collections). If the case is sent to ACS, ACS personnel will try to contact the taxpayer by correspondence and by phone to work with the taxpayer to find a payment solution. If ACS is unsuccessful at resolving the unpaid balance, the case is sent directly to a local IRS field office in which a Revenue Officer will work with directly with the taxpayer to attempt to resolve the balance due.

To calculate the costs of these functions, we use data from 2016 (the earliest year for which information on Collections is available) from the internal IRS accounting data for Notices, ACS, and Field Collections.

Per-collections notice. The cost multiplier we use for notices is a cost-per-notice multiplier. These values are highlighted in Table A.VIII. The average cost is \$10.97 for sending a first notice, \$9.13 for a second notice, and \$17.70 for a "final" notice.

We identify someone as having received a notice (for notices 1–4) if they have positive revenue associated with their audit from any of the parts of the collections process as marked in Table A.IX. We apply the average per-notice rate for "final" notices to the third and fourth notices.

appeals and slightly overestimating some other overhead items, but our estimate of the overall costs of the audit is largely unaffected.

Appendix Table A.VIII:
Average Cost per Notice ($)

Year	Average Cost per Notice		
	1st Notice	2nd Notice	Final Notice
2016	10.97	9.13	17.7
2017	13.22	10.25	16.32
2018	15.68	12.5	17.97
2019	12.7	12.6	16.3
2020	15.61	16.4	20.69

Appendix Table A.IX:
Per-Notice Average Rate Applied if Positive Revenue Found from These Stages

If positive revenue from:	Notice 1 Average Rate Applied	Notice 2 Average Rate Applied	Notice 3 Average Rate Applied	Notice 4 Average Rate Applied
Notice 1	X			
Notice 2	X	X		
Notice 3	X	X	X	
Notice 4	X	X	X	X
ACS	X	X	X	X
Collections queue	X	X	X	X
Field collections	X	X	X	X

ACS. We estimate the cost multiplier from an audit going through ACS as a "cost-per-dollar raised." We use the 2016 value of cost-per-dollar raised ($0.0513), as shown in Table A.X.

Field collections. We apply our estimated cost multiplier from in-person exams to estimate the total costs associated with field collections. This is because we are not able to separately calculate direct labor costs that result from time in field collections (as opposed to direct labor costs associated with the entire collections process). We assume that, on average, the ratio of non-direct labor costs to direct labor costs is the same between in-person exams and in-person collections. This assumption is simpler and, in our view, more plausible than the various assumptions needed for our best attempt to calculate a separate, field collections-specific cost multiplier (which yielded a cost multiplier of 2.19).

Appendix Table A.X:
ACS Multiplier (Cost per Dollar of Revenue)

Year	Cost per Revenue Dollar
2016	$\boxed{0.0513}$
2017	0.0551
2018	0.0567
2019	0.0529
2020	0.0518

C Correspondence Audits

In this section, we present results for the average returns to correspondence audits. Correspondence audits are primarily conducted by mail and are not assigned to a particular auditor. Correspondence audits are generally less complex than in-person audits and rely more heavily on algorithms to identify noncompliance. Appendix Figure A.II presents the cost per audit and revenue raised in each stage of the audit process, using the same format as Figure I for in-person audits. Labor costs are a much smaller fraction of the total costs of a correspondence audit than of an in-person audit. However, the overall ratio of revenue to cost for correspondence audits after accounting for non-labor costs is 2.13, similar to the ratio for in-person audits.

Appendix Figure A.III shows how these patterns vary across the income distribution. Panel A shows correspondence audit rates across the income distribution. The high rates of audits in the 20–30th income percentiles are EITC correspondence audits that typically ask for verification of relationship and residency status (e.g. forms such as birth certificates, school paperwork showing home address, etc.). Panel B shows how the hours spent and wages of auditors vary across the income distribution. The hours spent per audit are significantly lower for correspondence audits than for in-person audits. As with in-person audits, correspondence audits of higher-income taxpayers require higher-paid auditors and more of auditors' time. Panel C shows how revenues and costs per correspondence audit vary across the income distribution. Total costs per audit are broadly increasing in taxpayer income; but as with in-person audits, revenues increase faster than costs. Panel D shows the ratio of revenue to total costs per audit. This ratio rises from below 1 for audits at the bottom of the income distribution to 3.4 for audits of individuals in the 99–99.9th percentile of the TPI distribution. It then increases to 11.7 in the top 0.1% of the income distribution.

To summarize, non-labor costs comprise a much larger share of the costs of correspondence audits than of in-person audits, but the average return to correspondence audits is similar to the return to in-person audits, and returns to both types of audits increase with taxpayer income.

D Relation to Existing Estimates

In this appendix, we discuss the relationship of our estimates to two key estimates from existing literature on the returns to audits: Holtzblatt and McGuire [2020] and Sarin and Summers [2019].

Holtzblatt and McGuire [2020] (referred to as HM2020 hereafter) estimate the revenues and costs associated with IRS operational audits in the US. They estimate that, for in-person audits, revenue collected divided by costs was 3.3 and 2.8 in 2010 and 2017, respectively. It is not possible to conduct an exact comparison between that paper and this one because of differences between the two papers' samples. For example, HM2020 include not only individual audits, but also audits of corporations. They examine audits from tax years 2010 and 2017 that were completed before March 31st, 2012 and 2019. They exclude returns with EITC and also remove any "outlier" returns, defined as those in the top 0.5% of taxes collected.

One key difference that leads to higher average returns in HM2020 as compared to our work is that the cost estimates in HM2020 only include the direct costs of hours spent on enforcement activity and do not include additional labor and non-labor costs. In particular, HM2020 calculate the cost of direct enforcement activity by multiplying hours spent auditing by a wage rate. They then incorporate an estimate of additional employee benefit costs on top of those hourly wage costs. As they note, this does not include labor costs for management and support staff. It appears to omit the labor costs of time spent by enforcement personnel on non-audit tasks and does not include non-labor costs such as rent and IT costs. In our cost estimates, which are based on internal IRS accounting information, we find that non-labor costs contribute a substantial fraction of total costs. We estimate that, on average, total costs are 4.39 times larger than the direct labor cost of enforcement activities, and we estimate that more than half of that 4.39 figure is the result of non-labor costs.[14] The inclusion of these non-labor costs in our average cost estimates contributes to the divergence between our estimates and HM2020.[15]

Sarin and Summers [2019] discuss the returns to auditing very high income taxpayers and argue that auditing taxpayers with more than $5 million in earnings can produce a return of 18:1. In particular, they conduct a back-of-the-envelope calculation, drawing upon hourly audit adjustment estimates from George [2019] and dividing by estimates of average auditor costs. While the broad trajectory of our results are consistent with their findings, the approach used in our paper differs from the approach in theirs. First, George's estimates quantify recommended audit adjustments rather than audit revenue collected. A meaningful portion of assessed tax obligations are overturned on appeal and never collected, and so these adjustments often far exceed audit revenue. This means that the observed return on audit expenditures should fall below the ratio of audit assessments to auditor costs. Second, if the calculation in George [2019] is focused on the hours necessary to produce audit adjustments, these calculations may omit the cost of auditor hours following the initial recommendation and appeals and collections activity. While there is no formal confirmation of this hypothesis in George [2019], we find the average hours associated with an audit of the top 0.1% are meaningfully in excess of the average hours reported in that work.[16] Finally, our analysis of very high income taxpayers focuses on those in the top 0.1%. This threshold falls below the $5 million threshold in Sarin and Summers [2019] and so if the pattern of increasing retrurns with income continues to hold we should expect a lower return for the top 0.1% than for those with at least $5 million of income. Those differences help to explain the discrepancy between 18:1 figure

[14]HM2020 argue that non-labor costs are likely to be small because 94% of the IRS enforcement budget is attributable to personnel compensation. That calculation, based on Table 28 on the 2018 IRS Data Book, appears to only include costs associated with the "enforcement" line-item and, therefore, omits the costs associated with the "operational support" line-item. The lion's share of non-labor costs can be found under the operational support line item.

[15]Our marginal cost estimates in Section 7 fall closer to the estimates in HM2020 because, while they still include the labor costs associated with support staff and non-auditing hours, they do not include the portion of non-labor costs we estimate to be fixed.

[16]Hours per audit differ because of different definitions of audit, but the total revenue estimates in George (2019) resemble the revenue estimates in our work.

in Sarin and Summers [2019] and the 6:1 average return we find for the top 0.1% when deterrence effects are not included.[17]

Our estimates are comparable to recent budget scores produced by the Congressional Budget Office. For example, the CBO estimated (with no deterrence effects) that $46 billion in additional audit enforcement focused on high-income taxpayers would raise $180 billion in revenue, a revenue to cost ratio around 3.9. That 3.9 figure aligns fairly closely with our estimates of the returns to marginal audits of taxpayers of high income taxpayers absent deterrence effects. We find returns of 3.2:1 in the 99–99.9th percentile and 6.2:1 in the top 0.1%.

E MVPF of Tax Evasion

In this Appendix, we provide a class of structural models that motivate our MVPF formula in Equation (1). Our modeling approach builds on the large literature on tax evasion (e.g. Allingham et al. (1972); Keen and Slemrod (2017)), but extends to a dynamic context that allows for audits today to change tax payments and evasion behavior in the future. In order to incorporate dynamics while still keeping the model relatively tractable, we economize on other features of the model. For example, we assume quasilinear utility, we do not allow the probability of audits to depend on past behavior, and we do not allow for strategic interactions between evasion levels and probability of audits. We show that, with these assumptions, we can derive our exact formula for the MVPF in Equation (1). We then discuss how relaxing each of these assumptions leads to potential modifications to the MVPF. Finally, we discuss how one can use the MVPF combined with social welfare weights to compare the desirability of raising revenue through audits to other potential policies.

E.1 Setup

We consider an individual, i, who has a utility function over consumption, c_t, and earnings, y_t, in each period indexed by t. Earnings are taxed at $T(y_t)$ so that in the absence of any evasion consumption would be equal to $y_t - T(y_t)$. However, individuals have the opportunity to evade e_t dollars of their tax liability. This increases consumption by e_t in the event they are not audited. We let a_t denote an indicator for being audited in period t, and $\alpha_t = (a_1, ..., a_{t-1})$ denote the individual's audit history up through period t. We assume that, when audited, the individual must repay the evaded amount, e_t, plus a penalty $\phi^{\alpha_t}(e_t)$ that depends not only on the evaded amount but also on the individual's audit history, α_t.

Utility in each period is given by $u_i(c_t, y_t) = c_t - \psi_i(y_t) - \mathbb{I}\{a_t = 1\} B_i$, where $\psi_i(\circ)$ measures individual i's disutility of earning income and B_i measures their disutility of being audited (alternatively, the "taxpayer burden" of being audited). As noted above, we assume for simplicity that evasion has no psychic cost. This means the expected PDV of utility is given by

$$U_i = \mathrm{E}\left[\sum_{t=1}^{\infty} \beta^{t-1}\Big(c_t - \psi_i(y_t) - \mathbb{I}\{a_t = 1\} B_i\Big)\right]$$

where the expectation is taken with respect to the probability that $a_t = 1$ in each period, which we denote $p_t = \Pr\{a_t = 1 | e_t\}$. We assume for simplicity this probability is exogenous to evasion choices and income, and that the probability of future audit does not depend on past audits. As we discuss

[17] Coincidentally, when deterrence effects are included, the returns we find at the top of the income distribution approach or exceed 18:1.

further below in Section E.4.3, our formulas are unchanged when allowing audit probabilities to depend on evasion levels, $p_t(e_t)$.[18]

In each period, the budget constraint is given by

$$\begin{aligned} c_t &\leq y_t - T(y_t) + e_t \quad \text{if } a_t = 0 \\ c_t &\leq y_t - T(y_t) - \phi^{\alpha_t}(e_t) \quad \text{if } a_t = 1 \end{aligned}$$

We note that, while we model ϕ^{α_t} as the true penalty, it would be straightforward to extend the model to allow for misperceptions of the penalty by interpreting ϕ^{α_t} as the perceived penalty.

Individuals make two choices: earnings and evasion. The additive separability in the model implies that the choice of earnings y_t is independent of audits and/or evasion: The optimal choice of income, y_t^*, satisfies $\psi_i'(y_t^*) = 1 - T'(y_t^*)$ in each period (we suppress the i subscript on y_t^* but note this choice differs across individuals). After plugging the budget constraints into the objective function, we see that the choice of evasion in each period solves:

$$\max_{e_t} (1 - p_t(e_t)) e_t - p_t(e_t) \phi^{\alpha_t}(e_t) \tag{1}$$

Intuitively, individuals maximize the expected money they keep from the government net of penalties. In other words, they minimize the expected taxes they pay inclusive of expected penalties that they pay in the event they are audited. We let e_{t,α_t}^* denote the solution to this maximization program in each period t after realizing audit history, α_t. This is given by:

$$e_{t,\alpha_t}^* = \left(\frac{\partial \phi^{\alpha_t}}{\partial e}\right)^{-1} \left(\frac{1 - p_t}{p_t}\right)$$

Note that this equation shows how audits can impact future evasion behavior: if being audited increases the marginal penalty from future evasion (e.g. because it is no longer possible to claim that misreporting was not willful), individuals may choose to reduce their evasion behavior in the future.

We can then plug in the choice of evasion into the utility function to write the indirect expected ex ante utility as:

$$\begin{aligned} V_i(\{p_t\}, \{\phi^{\alpha_t}(\circ)\}, T(\circ)) &= \mathrm{E}\left[\sum_{t=1}^{\infty} \beta^{t-1} \left(y_t^* - T(y_t^*) - \psi_i(y_t^*) + e_{t,\alpha_t}^* - \mathbb{I}\{a_t = 1\} \left(e_{t,\alpha_t}^* + \phi^{\alpha_t}\left(e_{t,\alpha_t}^*\right) + B_i\right)\right)\right] \\ &= \sum_{t=1}^{\infty} \beta^{t-1} \Pr\{\alpha_t\} \left(y_t^* - T(y_t^*) - \psi_i(y_t^*) + e_{t,\alpha_t}^* - p_t\left(e_{t,\alpha_t}^* + \phi^{\alpha_t}\left(e_{t,\alpha_t}^*\right) + B_i\right)\right) \end{aligned}$$

where $\Pr\{\alpha_t\}$ is the probability of a particular audit sequence, α_t. The *ex ante* expected utility experienced by the individual, V, is a function of the audit probabilities, p_t, and the penalty functions, $\phi^{\alpha_t}(\circ)$, and the tax schedule, $T(\circ)$.

E.2 Willingness to Pay for Expanded Audits

We now can ask: what is the welfare impact of expanding audits? We model this as an increase in the audit probability in the first period by dp_1. Individuals are willing to pay $\frac{dV}{dp_1}$ in order to avoid an audit. To see how changing p_1 affects V_i, it is helpful to write V_i by expanding out the first period probability of audit. We have:

[18] In Section E.4.3, we also discuss a small modification below that enables the probability of a future audit to be increasing in the presence of past audits.

$$V_i = p_1 \left[y_1^* - T\left(y_1^*\right) - \psi\left(y_1^*\right) - \phi^{\alpha_1}\left(e_{1,\alpha_1}^*\right) - B + \beta V_i^1 \right] + (1-p_1)\left[y_1^* - T\left(y_1^*\right) - \psi\left(y_1^*\right) + e_{1,\alpha_1}^* + \beta V_i^0 \right]$$

where $\phi^{\alpha_1}(\circ)$ is the penalty in the first period (before there is any audit history for the individual) and e_{1,α_1}^* is the choice of evasion in the first period. The first term in brackets is the utility if audited and the second term is the utility if not audited. The term V_i^1 is the PDV of future utility in subsequent periods if $a_1 = 1$ and V_i^0 is the PDV of future utility in subsequent periods if $a_1 = 0$. The envelope theorem implies that the impact of increasing p_1 affects utility through both the first period utility and the impact on future utility:

$$\begin{aligned} -\frac{dV_i}{dp_1} &= e_1^{a_t} + \phi^{\alpha_1}\left(e_{1,\alpha_1}^*\right) + B_i + \left(V_i^0 - V_i^1\right) \\ &= R^{mech} + B + \left(V_i^0 - V_i^1\right) \end{aligned}$$

The first period utility impact of the additional audits is given by the level of evasion plus the penalty and the taxpayer burden. To calculate the impact of the audit on future periods, note that we can write $V_i^0 - V_i^1$ as the present discounted future revenue collected by the government from reduced evasion:

$$\begin{aligned} V_i^0 - V_i^1 &= \sum_{t=2}^{\infty} \beta^{t-1} \left(E\left[e_{t,\alpha_t}^* - p_t\left(e_{t,\alpha_t}^* + \phi^{\alpha_t}\left(e_{t,\alpha_t}^*\right)\right) | a_1 = 0 \right] - E\left[e_{t,\alpha_t}^* - p_t\left(e_{t,\alpha_t}^* + \phi^{\alpha_t}\left(e_{t,\alpha_t}^*\right)\right) | a_1 = 1 \right] \right) \\ &\equiv R_i^{future} \end{aligned}$$

so that $V_i^0 - V_i^1$ is the causal effect of the audit in period 1 on the PDV of the change in tax revenue paid to the government in the future as a result of being in the audited ($\alpha_1 = 1$) versus non-audited ($\alpha_1 = 0$) state of the world in period 1. Combining, the willingness to pay to avoid an expansion of audits in period 1 is given by:

$$-\frac{dV}{dp_1} = R_i^{mech} + R_i^{future} + B_i \tag{2}$$

which is the sum of the mechanical revenue collected by the audit, the future PDV revenue collected as a result of within-person deterrence from the audit, and the taxpayer burden of the audit.[19]

Equation (2) provides a measurement of a given individual's willingness to pay to avoid the audit. Those with higher financial impacts or non-financial burden from the audit have a higher willingness to pay to avoid the audit. With a slight abuse of notation, we let R^{mech},R^{future},and B without i subscripts denote the average values of these variables among those being audited.

E.3 Government Revenue and the MVPF of Expanded Audits

Let G denote the PDV of government revenue:

$$G = \sum_{t=1}^{\infty} \beta^{t-1} \mathrm{E}\left[T\left(y_t^*\right) - e_{t,\alpha_t}^* + \mathbb{I}\left\{a_t = 1\right\}\left[e_{t,\alpha_t}^* + \phi^{\alpha_t}\left(e_{t,\alpha_t}^*\right) - C \right] - F \right]$$

[19] We assume that the probability of audit p_t is independent of the choice of income that an individual has. This latter assumption can easily be relaxed by assuming that income choices, y_t, are affected by the probability of the audit. The envelope theorem implies that these will not enter the willingness to pay to avoid the expanded audits. They could, however, generate an additional positive or negative revenue to the government from the audit.

where C is the marginal cost of an audit and F is the fixed costs of audits. We assume the government and individuals have the same expectations and discount factor.

The effect of expanding audits in period 1 is given by the sum of the revenue collected in the first period, R^{mech} and the revenue collected in future periods R^{future} minus the marginal cost of the audits in period 1:

$$\frac{dG}{dp_1} = R^{mech} + R^{future} - C \tag{3}$$

Let $R = R^{mech} + R^{future}$ denote the total PDV of government revenue collected as a result of the audit. Combining the willingness to pay by audited individuals to avoid an audit and the revenue raised by audits, the MVPF of individual tax audits can be expressed as:

$$MVPF = \frac{R+B}{R-C}$$

which is precisely our formula in equation (1).

Specific versus General Deterrence This MVPF expression does not contain a term for the deterrence effect of the increased probability, p_t, on evasion in the first period, e^*_{1,α_1}. The envelope theorem implies that the change in evasion in response to the increase in threat of audit in period 1 does not enter the willingness to pay term. For the government cost, note that the net impact on the government budget from changes in e^*_{1,α_1} is given by:

$$-\frac{de^*_{1,\alpha_1}}{dp_t}\left[(1-p_1) - p_1\frac{\partial \phi^{\alpha_1}}{\partial e}\left(e^*_{1,\alpha_1}\right)\right] = 0 \tag{4}$$

Because individuals maximize their expected income from evasion, they also minimize the government revenue from evasion. Recall from the individuals' optimization problem above that $\frac{\partial \phi^{\alpha_t}}{\partial e}\left(e^*_{t,\alpha_t}\right) = \frac{1-p_t}{p_t}$. This means that small changes in evasion behavior in response to changes in the audit probability do not affect government revenue. Expanding audits leads people to evade less ($\frac{de}{dp} < 0$), but the loss in penalties perfectly offsets the gain in tax revenue from reduced evasion. As a result, general deterrence (i.e. $\frac{de^*_{1,\alpha_1}}{dp_t}$) does not enter either the numerator or denominator of the MVPF. We show below that, under alternative assumptions about taxpayer behavior, one can obtain general deterrence effects that change government revenue. In these cases, general deterrence should be included in the net cost (denominator) of the MVPF but the envelope theorem continues to imply that it does not enter the numerator.

At this point, one might be puzzled why the willingness to pay to avoid an audit includes individual deterrence, R^{future}, but not a term related to general deterrence, $-\frac{de^*_{1,\alpha_1}}{dp_t}$. We noted above that the envelope theorem applies to $-\frac{de^*_{1,\alpha_1}}{dp_t}$, so why doesn't the envelope theorem also apply to individual deterrence? The key distinction is that the individual deterrence term includes not only a behavioral response but also the mechanical revenue collected from changes to the future constraints faced by an audited taxpayer. When individuals are audited, their future audit penalties, ϕ^{α_t}, from continuing their evasion increase.[20] The impact on revenue, R^{future}, is then the sum of the mechanical and behavioral components from these increased penalties. To see this, consider just the revenue collected in period 2 as a result of expanded audits in period 1. Let R_2^{future} denote this period 2 revenue and let e_2^0 and e_2^1 denote the amount evaded in period $t = 2$ conditional on

[20] See Section E.4.3 below for a discussion of the case when probabilities of future audit increase.

having been not audited or audited in period 1. The impact of the expanded audit in period 1 on revenue in period 2 is given by:

$$R_2^{future} = e_2^0 - p_2\left(e_2^0 + \phi^0\left(e_2^0\right)\right) - \left(e_2^1 - p_2\left(e_2^1 + \phi^1\left(e_2^1\right)\right)\right) \tag{5}$$

Now, note that we can write the change in penalty revenue as:

$$\phi^0\left(e_2^0\right) - \phi^1\left(e_2^1\right) = \phi^0\left(e_2^0\right) - \phi^1\left(e_2^0\right) + \phi^1\left(e_2^0\right) - \phi^1\left(e_2^1\right)$$

where $\phi^0\left(e_2^0\right) - \phi^1\left(e_2^0\right)$ is the mechanical effect on audit revenue of increased penalties after an audit holding evasion in period 2 fixed and $\phi^1\left(e_2^0\right) - \phi^1\left(e_2^1\right)$ is the impact of the change in evasion on the penalty revenue. Plugging this back into R_2^{future} we obtain

$$R_2^{future} = \underbrace{\phi^0\left(e_2^0\right) - \phi^1\left(e_2^1\right)}_{Mechanical} + \underbrace{\left(e_2^0 - e_2^1\right)\left(1 - p_2 - \frac{\phi^1\left(e_2^0\right) - \phi^1\left(e_2^1\right)}{e_2^0 - e_2^1}\right)}_{Behavioral}.$$

This equation shows that the impact of an additional audit in period 1 on revenue in period 2 is the sum of two components: the mechanical revenue obtained from higher penalties at the level of evasion chosen by someone who had not been audited and the impact of the change in evasion, $e_2^0 - e_2^1$,

In contrast, the impact of general deterrence on government revenue is given by Equation (4) above, which equals zero due to the envelope theorem from individuals maximizing utility when choosing their evasion. Intuitively, the general deterrence term is the behavioral component that comes along with the mechanical revenue component given by the revenue collected by the audits in period 1. The individual deterrence effect is, in contrast, the sum of both the mechancial and behavioral revenue components in future periods.

There are therefore two differences between specific and general deterrence. First, individual deterrence contains a mechanical revenue effect due to the higher penalties faced by taxpayers if they are caught misreporting again. In contrast, general deterrence does not have a mechanical revenue component: it is defined as solely the behavioral response component to the increase in p. Second, even though the change in audits is marginal, an audit leads to a non-marginal change in the future constraints faced by the individual (i.e. strictly higher penalties). The threat of future penalties may lead taxpayers to the corner solution where they report truthfully. Because these increases in penalties can be large, it is important when measuring individual deterrence to include the impact of the behavioral response to the policy on the government budget. In our model, the quasilinearity of preferences means that we can instead measure the welfare cost (mechanical + behavioral components) from individual deterrence as the net impact on future tax revenue.

E.4 Robustness and Extensions

The baseline model above makes a series of modeling assumptions. In this subsection, we discuss a few of these key assumptions and how relaxing them changes the MVPF. We begin by showing how one can use a Baily-Chetty style adjustment to our formulas to account for the fact that individuals may be risk averse when facing audits. Next, we discuss issues related to general deterrence and how various forms of general deterrence can enter the MVPF. Finally, we discuss the role of endogenous audit probabilities.

E.4.1 Risk Aversion

Let $u(c)$ denote utility over realized consumption. In the presence of risk aversion, V_i now becomes:

$$V_i = p_1 \left[u\left(y_1^* - T(y_1^*) - \psi(y_1^*) - \phi^{\alpha_1}\left(e_{1,\alpha_1}^*\right) - B\right) + \beta V_i^1\right] + (1-p_1)\left[u\left(y_1^* - T(y_1^*) - \psi(y_1^*) + e_{1,\alpha_1}^*\right) + \beta V_i^0\right]$$

So, when we consider the derivative with respect to p_1 we end up with

$$-\frac{dV_i}{dp_1} = u(c_{noaudit}) - u(c_{audit}) + \left(V_i^0 - V_i^1\right)$$

where

$$\begin{aligned} c_{audit} &= y_1^* - T(y_1^*) - \psi(y_1^*) - \phi^{\alpha_1}\left(e_{1,\alpha_1}^*\right) - B \\ c_{noaudit} &= y_1^* - T(y_1^*) - \psi(y_1^*) + e_{1,\alpha_1}^* \end{aligned}$$

Following Baily (1978), we can take a second order Taylor expansion of the utility function around a fixed level of consumption to approximate the difference in the levels of utility with the marginal utility of consumption yielding

$$u(c) \approx u(c_0) + u'(c_0)(c - c_0) + \frac{1}{2}u''(c_0)(c - c_0)^2$$

where u' and u'' are the values of the first and second derivative of the utility function at a particular point of consumption. We take the point of marginal utility expansion to be the one normalizing willingness to pay on the left-hand side of the equation. This leaves us with

$$\begin{aligned} u(c_{audit}) - u(c_{noaudit}) &\approx u'(c_{noaudit})\Delta c - \frac{1}{2}u''(c_{noaudit})\left[(c_{audit} - c_{noaudit})^2\right] \\ &= u'(c_{noaudit})\Delta c\left(1 + \frac{1}{2}\sigma\frac{\Delta c}{c}\right) \end{aligned}$$

where $\sigma = \frac{-u''(c_{noaudit})c_{noaudit}}{u'(c_{noaudit})}$ is the coefficient of relative risk aversion evaluated at the level of no audit consumption and $\frac{\Delta c}{c} = \frac{c_{audit} - c_{noaudit}}{c_{audit}}$ is the change in consumption relative to the consumption in the absence of an audit. This equation shows that willingness to pay to avoid the audit is higher than the full monetary difference in consumption between audited and non-audited states of the world. We note that a similar adjustment applies to $V_i^0 - V_i^1$, so that one would also adjust upwards the willingness to pay from individual deterrence.

The advantage of this expression is that it provides guidance on the magnitude of how risk aversion affects the MVPF (recall the net cost to the government is unchanged): the willingness to pay is adjusted upwards by the coefficient of relative risk aversion times the percent change in consumption as a result of the audit. The ideal dataset would track consumption of audited versus non-audited individuals, in the same way such formulas have been applied in the unemployment insurance context. To start, consider the top of the income distribution. Audits of taxpayers in the top 10% of the income distribution collect revenue around \$30,000, and we know that roughly a third of that amount is the persistent increase in reported tax liability. Hence, it would be natural to think of audits as leading to a drop in consumption of \$10,000, which is at most 5% of average income. Combining this with a CRRA of 3 would suggest a willingness to pay that is 15% higher to avoid the audit. This would increase the MVPF from 1.15 to 1.3. Even under these conservative assumptions, the MVPF of expanding audits of top earners is less than the estimated MVPF of 1.5 or more of increasing tax rates on the same group.

We can consider a similar exercise at the bottom of the income distribution. Here, perhaps the upfront revenue paid of \$5,000 translates into a consumption drop of around \$2,000, which would be a roughly 10% drop for someone consuming \$20,000. With a coefficient of relative risk aversion of 3, this would suggest we need to revise upwards the MVPF for audits of low income taxpayers by 30%. While future work should seek to measure the consumption impact of audits across the income distribution, these stylized calculations suggest that including risk aversion would not overturn our core conclusion that the MVPF of audits is declining across the income distribution.

E.4.2 General Deterrence

Our baseline formula does not include general deterrence effects. This is because when individuals choose evasion levels to maximize income after taxes and expected penalties, they also minimize government revenue from taxes and penalties so that the envelope theorem from taxpayer maximization also implies that general deterrence responses from audits have no first order effect on government revenue.

In this section, we discuss and relax the modeling assumptions that generate this result. We show that this result is dependent on the combination of (a) no risk aversion and (b) no "non-financial" costs of evasion. When either of these assumptions does not hold, then changes in evasion in response to the increased threat of audit can have first order fiscal benefits/costs to the government. These fiscal effects should be included in the denominator of the MVPF. However, they continue to be omitted from the numerator of the MVPF due to the envelope theorem. This means that including general deterrence effects is likely to reduce the MVPF of tax audits, which makes expanded audits a more efficient method of raising revenue than is implied by our baseline measures.

We discuss the cases of risk aversion and non-financial costs of evasion in turn below.

Risk Aversion Suppose individuals are risk averse. Let $u'_{1,0}$ be the marginal utility of consumption in period 1 if not audited and $u'_{1,1}$ be the marginal utility of consumption in period 1 if audited. Note that, as in Allingham et al. [1972], we expect $u'_{1,1} > u'_{1,0}$ because if caught the individual must repay the evaded amount plus penalties. Individuals choose evasion in period 1, e_{1,α_1}, to maximize expected utility, leading to the first order condition:

$$-\frac{de^*_{1,\alpha_1}}{dp_t}\left[(1-p_1)\,u'_{1,0} - p_1\frac{\partial\phi^{\alpha_1}}{\partial e}\left(e^*_{1,\alpha_1}\right)u'_{1,1}\right] = 0$$

which can be re-arranged to yield:

$$\left[(1-p_1) - p_1\frac{\partial\phi^{\alpha_1}}{\partial e}\left(e^*_{1,\alpha_1}\right)\right] \quad = \quad p_1\frac{\partial\phi^{\alpha_1}}{\partial e}\left(e^*_{1,\alpha_1}\right)\left(\frac{u'_{1,1}}{u'_{1,0}} - 1\right)$$

With this re-arrangement, the left-hand side is the lost government revenue from additional evasion and the right-hand side is the wedge from risk aversion (which is positive as long as the audit leads to an increase in the marginal utility of consumption). This is similar to Equation (4) but also accounts for higher marginal utility of income in the audited state of the world. As a result, a marginal increase in evasion has a first order effect on government revenue. In the model, this general deterrence effect on government revenue is given by $-\frac{de^*_{1,\alpha_1}}{dp_t}\left[(1-p_1) - p_1\frac{\partial\phi^{\alpha_1}}{\partial e}\left(e^*_{1,\alpha_1}\right)\right]$ or equivalently by $-\frac{de^*_{1,\alpha_1}}{dp_t}p_1\frac{\partial\phi^{\alpha_1}}{\partial e}\left(e^*_{1,\alpha_1}\right)\left(\frac{u'_{1,1}}{u'_{1,0}} - 1\right)$. This additional term would enter as an additional fiscal externality in the denominator of the MVPF. However, the change in evasion in response to the increased probability of audit would not enter the individual's willingness to pay due to the envelope theorem. In this specification, general deterrence effects would lead to a lower MVPF than in our baseline estimate that does not allow for risk aversion, making audits a more welfare efficient method of raising revenue.

Non-Financial Costs Next, suppose that we have our baseline model but now there are non-financial costs $C(e_t)$ of evasion that are increasing in the size of evasion, $C'(e_t) > 0$. These could include either a moral cost to individuals from dishonesty or other non-financial penalties such as jail time. The individual now chooses evasion to maximize:

$$\max_{e_t} (1-p_t)\, e_t - p_t \phi^{\alpha_t}(e_t) - C(e_t)$$

which generates a first order condition for the optimal choice e^*_{1,α_1}:

$$-\frac{de^*_{1,\alpha_1}}{dp_t}\left[(1-p_1) - p_1 \frac{\partial \phi^{\alpha_1}}{\partial e}\left(e^*_{1,\alpha_1}\right) - C'\left(e^*_{1,\alpha_1}\right)\right] = 0$$

which can be re-written as:

$$\left[(1-p_1) - p_1 \frac{\partial \phi^{\alpha_1}}{\partial e}\left(e^*_{1,\alpha_1}\right)\right] \quad = \quad C'\left(e^*_{1,\alpha_1}\right)$$

Note that the left-hand side is the marginal impact of general deterrence, $-\frac{de^*_{1,\alpha_1}}{dp_t}$, on government revenue, as in equation (4), and the right-hand side is the slope of the non-financial costs with respect to the size of the evasion. As with risk aversion, a reduction in evasion raises government revenue and should be incorporated into the denominator of the MVPF. Here again, however, the general deterrence response, $\frac{de^*_{1,\alpha_1}}{dp_t}$, does not affect the individual's willingness to pay to avoid the audit due to the envelope theorem.

To summarize, under the benchmark case of risk neutrality and purely financial costs from audits, general deterrence effects have no impacts on government revenue. When either of these assumptions is relaxed, general deterrence responses can have impacts on government revenue. Future work should aim to estimate these general deterrence effects and include them in the denominator (but not the numerator) of the MVPF.

E.4.3 Endogenous Audit Probability

Our baseline parameterization assumes that the probability of an audit is exogenous and independent of past audits. Here, we relax both of these assumptions, allowing the probability of an audit to depend on the extent of past audits and the amount evaded. We begin with the case where past audits increase the probability of future audits.

Past Audits Increase Probability of Future Audits Our baseline specification assumes the core reason that people reduce evasion after being audited is that they fear higher penalties in the future. This modeling choice is motivated by the sharply increasing penalties for repeated and willful (as opposed to unintentional) misreporting. Nonetheless, an alternative potential reason that we find significant individual deterrence effects is that people anticipate a higher future audit probability after being audited (e.g. they expect that that are being watched more closely). Here, we assess how our MVPF calculations would differ if the individual deterrence effect is driven by changes in future audit probabilities as opposed to changes in future penalties. The key insight from this specification is that one needs to not only include the additional revenue from individual deterrence but also the greater future audit burdens and higher administrative costs to perform the additional audits that arise due to the higher future audit probabilities. Under the assumption that the impact of audit probabilities on evasion is perfectly linear, the MVPF reduces to the static MVPF excluding individual deterrence. This is because any future revenue gained from changes in evasion (individual deterrence) is offset by reduced revenue from future audits of that individual.

To allow future audit probabilities to depend on past audits, let $p_t(\alpha_t)$ denote the audit probability in period t given audit history α_t. Let $p_t^1(\alpha_t)$ denote the probability of a future audit in period t conditional on an audit in period 1 and $p_t^0(\alpha_t)$ denote the probability of a future audit in period t conditional on no audit in period 1. The willingness to pay to avoid the audit now has a similar but slightly different form (abstracting from individual heterogeneity by dropping i subscripts):

$$-\frac{dV}{dp_1} = R^{mech} + \left(V^0 - V^1\right) + B\left(1 + \sum_{t>1} \beta^{t-1} E\left[p_t^1(\alpha_t) - p_t^0(\alpha_t)\right]\right) \tag{6}$$

conditional where the burden term, B, is now multiplied by the increase in the number of future audits so that we capture the PDV of burdens experienced by an audited taxpayer inclusive of greater future audits. Here, the term $V^0 - V^1$ is again given by the future revenue collected from audits:

$$\begin{aligned} V_i^0 - V_i^1 &= \sum_{t=2}^{\infty} \beta^{t-1} \left(E\left[e_{t,\alpha_t}^* - p_t^0(\alpha_t)\left(e_{t,\alpha_t}^* + \phi^{\alpha_t}\left(e_{t,\alpha_t}^*\right)\right)|a_1 = 0\right] - E\left[e_{t,\alpha_t}^* - p_t^1(\alpha_t)\left(e_{t,\alpha_t}^* + \phi^{\alpha_t}\left(e_{t,\alpha_t}^*\right)\right)|a_1 = 1\right]\right) \\ &\equiv R_i^{future} \end{aligned}$$

which now includes the revenue collected from the higher rate of future audits and higher penalties from future audits. To isolate the role of the probability of future audits, it is helpful to assume that the penalties do not vary with past audits (we note this is inconsistent with IRS practice, but helpful for understanding the general properties of the MVPF of audits). Consider again the future revenue to the government in period 2. Let p_2^0 denote the probability of being audited in period 2 if the individual is not audited in period 1 and p_2^1 the corresponding probability for those audited in period 1. The impact of the period 1 audit on revenue collected in period 2 (as measured in equation (5) above) now becomes

$$R_2^{future} = e_2^0 - p_2^0\left(e_2^0 + \phi\left(e_2^0\right)\right) - \left(e_2^1 - p_2^1\left(e_2^1 + \phi\left(e_2^1\right)\right)\right)$$

where we remove superscripts for ϕ to denote the fact it does not depend on past audits. Rearranging,

$$R_2^{future} = e_2^0 - e_2^1 - \left(p_2^0 - p_2^1\right)\left(e_2^0 + \phi\left(e_2^0\right)\right) + p_2^1\left(e_2^1 + \phi\left(e_2^1\right) - e_2^0 + \phi\left(e_2^0\right)\right)$$

so that the revenue collected in period 2 normalized by the change in audit probabilities is given by

$$\frac{R_2^{future}}{p_2^1 - p_2^0} = e_2^0 + \phi\left(e_2^0\right) + \frac{e_2^1 - e_2^0}{p_2^1 - p_2^0}\left(1 - p_2^1 \frac{e_2^1 + \phi\left(e_2^1\right) - e_2^0 + \phi\left(e_2^0\right)}{e_2^1 - e_2^0}\right)$$

The revenue collected in period 2 is the sum of the mechanical revenue, $e_2^0 + \phi\left(e_2^0\right)$, plus the impact of the behavioral response, $e_2^1 - e_2^0$, relative to the increased audit probabilities, $p_2^1 - p_2^0$, on net government revenue. To first order (i.e. $p_2^1 \approx p_2^0$), this latter term is zero because the choice of evasion minimizes government revenue — any gains from evasion, e_2, are weighed against the penalties from evasion $p_2\left(e_2 + \phi\left(e_2\right)\right)$, so that $e_2^1 \approx p_2^1\left(e_2^1 + \phi\left(e_2^1\right)\right)$ and $e_2^0 \approx p_2^1\left(e_2^0 + \phi\left(e_2^0\right)\right)$. To first order , the revenue from future audits is simply the mechanical revenue generated from the greater number of audits in period 2, $e_2^0 + \phi\left(e_2^0\right)$. Any decrease in evasion in period 2, e_2^0 leads to an offsetting reduction in revenue that is collected euring the audit.

If we now make the additional assumption that evasion is constant over time, $e_t^0 = e_1^0$, then $R^{mech} = \frac{R_t^{future}}{p_t^1 - p_t^0}$. Hence,

$$\begin{aligned} R^{future} = V^0 - V^1 &= \sum_{t>1} \beta^{t-1} R_t^{future} \\ &= R^{mech} \sum_{t>1} \beta^{t-1} E\left[p_t^1\left(\alpha_t\right) - p_t^0\left(\alpha_t\right)\right] \end{aligned}$$

which will become useful below when forming the MVPF.

Turning to government costs, an increase in audits today leads to more audits conducted in the future due to the dependence of p_t on past audits, α_t. The total cost is

$$\frac{dG}{dp_1} = R^{mech} + R^{future} - C\left(1 + \sum_{t>1} \beta^{t-1} E\left[p_t^1\left(\alpha_t\right) - p_t^0\left(\alpha_t\right)\right]\right) \tag{7}$$

which includes the additional costs from the greater probability of future audits. Combining the willingness to pay and net costs and using the fact that can divide through by $1+\sum_{t>1} \beta^{t-1} E\left[p_t^1\left(\alpha_t\right) - p_t^0\left(\alpha_t\right)\right]$, we arrive at an MVPF formula given by:

$$MVPF = \frac{R^{mech} + B}{R^{mech} - C}$$

This shows that to first order the MVPF only depends on the mechanical revenue generated by the audit. The individual deterrence effect increases future revenue but it is future audits generating that revenue. When evasion is constant over time, this means that the future revenue comes with the same proportional increase in government costs from those future audits. Hence, the key lesson here is that if individual deterrence is driven by the increase in future audit probabilities, then we need to account for the cost of those future audits when also measuring the benefits in terms of collected revenue. In contrast, if individual deterrence is driven by the threat of higher penalties, then the individual deterrence revenue is collected without additional cost to the government (which is related to the classic insight of Becker [1968] and Allingham et al. [1972] that higher penalties are more efficient than expanded audits, a point we return to in Section E.7 below). It is unlikely, however, that in our setting the deterrence effect is driven by increased audit probabilities. For our individual deterrence effects to be driven solely by increased audit probabilities, it would need to be the case that audited taxpayers face on average 3 additional audits in the 14 years after their initial audit; in practice, they face on average less than 1 additional audit. In contrast, in practice taxpayers can face serious financial and criminal penalties for repeated noncompliance.

Evasion Affects Audit Probabilities The previous analyses continued to maintain the assumption that the audit probability is independent of the level of evasion. Here, we relax this assumption and show that this largely does not affect the results. As in the previous case, we need to accurately measure the extent to which an audit today leads to more future audits so that we accurately measure audit burdens and costs; but if individual deterrence effects are driven by the increase in the slope of the audit probability with respect to evasion and do not lead to more future audits, our baseline formula is valid for measuring the MVPF of expanded audits.

To see this, we assume that evasion in period, e_t, increases the chance of being audited in period t. To capture this in a fairly general way, we let the audit probability depend on both the history of past audits and the current level of evasion, $p_t^{\alpha_t}(e_t)$, and assume this is continuously differentiable and increasing in evaded income. Given any audit history α_t, the generalization of equation (1) shows that optimal level of evasion e_{t,α_t}^* is given by the first order condition :

$$1-p_t^{\alpha_t}\left(e_{t,\alpha_t}^*\right) = \frac{\partial p_t^{\alpha_t}\left(e_{t,\alpha_t}^*\right)}{\partial e_t}\left(e_{t,\alpha_t}^*+\phi^{\alpha_t}\left(e_{t,\alpha_t}^*\right)\right)+p_t^{\alpha_t}\left(e_{t,\alpha_t}^*\right)\frac{\partial \phi^{\alpha_t}}{\partial e}\left(e_{t,\alpha_t}^*\right) \tag{8}$$

Evasion has benefits proportional to the probability of not being caught, $1-p_t^{\alpha_t}\left(e_{t,\alpha_t}^*\right)$, but has costs in terms of higher probability of being caught, $\frac{\partial p_t^{\alpha_t}\left(e_{t,\alpha_t}^*\right)}{\partial e_t}$ times the cost of being caught, $e_{t,\alpha_t}^*+\phi^{\alpha_t}\left(e_{t,\alpha_t}^*\right)$, plus the impact of the higher evasion on the fines paid, $p_t^{\alpha_t}\left(e_{t,\alpha_t}^*\right)\frac{\partial \phi^{\alpha_t}}{\partial e}\left(e_{t,\alpha_t}^*\right)$. With this formulation, individual deterrence can occur both because of higher future fines from evasion and from higher future audit probabilities. As noted in the previous section, these future burden and government costs need to be included if current audits lead to more future audits. However, aside from this modification, all of the formulas for the willingness to pay for the audit and cost to the government continue to hold when simply replacing p_t with the equilibrium level of $p_t^{\alpha_t}\left(e_{t,\alpha_t}^*\right)$.[21] For example, suppose the individual deterrence effect arises not because of higher future fines but rather because the IRS can more easily audit someone again and uncover their evasion — in other words, suppose $\frac{\partial p}{\partial e}$ is higher in future periods after an audit. This can lead to a reduction in evasion in future periods (i.e. is a rationale for individual deterrence). It need not lead to an increase in the actual number of future audits, but if it does this would need to be incorporated into the MVPF formula as in Equations (7) and (6).

Income Affects Audit Probability Lastly, suppose income choices affect audit probabilities. To do so, let the audit probability be given by $p(y;\mu)$ where y is the individual's choice of income and μ is a variational parameter describing potential changes to the audit function. For example, if $p_1(y)$ is the status quo audit function, we could write $p_1(y;\nu)=p_1(y)+\mu h(y)$ for a continuously differentiable positive function, h, and consider a small change in μ starting at $\mu=0$ (assuming the range of $p(y)$ is strictly in the interior of $(0,1)$). Formally, let $y_{i,1}^*(\mu)$ and $e_{i,1}^*(\mu)$ denote an individual's choices of income and evasion when faced with the probability of audit given by variation ν. We note that the willingness to pay to avoid the change in audit policy, $d\mu$, is given by:

$$-\frac{dV_i}{d\mu} = \frac{\partial p_1\left(y_1^*,e_1^*\right)}{\partial \mu}\left[e_1^{a_t}+\phi^{\alpha_1}\left(e_{1,\alpha_1}^*\right)+B_i+\left(V_i^0-V_i^1\right)\right]$$

so that

$$\frac{-\frac{dV_i}{d\mu}}{\frac{\partial p_1\left(y_1^*,e_1^*\right)}{\partial \mu}} = R^{mech}+B+\left(V_i^0-V_i^1\right)$$

In other words, the willingness to pay to avoid an additional audit is unchanged when we allow audit probabilities to be a function of income. This is a standard consequence of the envelope theorem.

However, the net cost to the government of a change in audit policy now includes fiscal externalities from changes in income in response to changes in the audit policy. It is straightforward to see that equation (3) now not only includes the mechanical and future revenue associated with the audit, but also includes the change in income choices, $\frac{\partial y}{\partial \mu}$, weighted by the marginal tax rate, $T'\left(y_1^*\right)$. For example, with only a single individual, we have

[21] In order for dV/dp_1 to be well-defined, one needs to add the natural assumption that the audit expansion increases p_1 for all e_1 (otherwise it wouldn't be clear what is meant by dp_1). This assumption is natural as we imagine expanded audits but not necessarily improvements in the audit technology (i.e. improvements in the slope of $p_t(e_t)$).

$$\frac{dG}{d\mu} = \frac{\frac{\partial y_1}{\partial \nu}}{\frac{\partial p_1}{\partial \nu}} T'\left(y_1^*\right) + R^{mech} + R^{future} - C$$

where $\frac{\frac{\partial y_1}{\partial \mu}}{\frac{\partial p}{\partial \mu}}$ is the impact of the policy on income choices, $\frac{\partial y_1}{\partial \mu}$, normalized by the net change in audit probability, $\frac{\partial p_1}{\partial \mu}$. If increasing audits on the rich causes people to no longer be rich, we need to estimate this behavioral response and multiply this change in behavior by the marginal tax rate on income, $T'\left(y_1^*\right)$. Hence, the MVPF of audits on the rich could have a higher MVPF when incorporating this feature.

Summary In sum, our core MVPF derivation relies on some assumptions — namely no risk aversion, no non-financial costs of the audit, and exogenous audit probabilities —- that can be relaxed if one is able to estimate the additional relevant empirical parameters. The presence of risk aversion would increase the MVPF (relative to what we measure) by a proportion equal to the percentage impact of the audit on consumption scaled by the coefficient of relative risk aversion. The general deterrence effect would be measured as the causal effect on government revenue of an additional audit due to changes in the *ex ante* evasion behavior of people regardless of whether they are audited (de_1/dp_1 in our model). The fiscal externality from these responses are zero under our baseline assumptions, but are nonzero in the presence of risk aversion or non-financial costs of audits. In this case, expanded audits could create general deterrence benefits to the government that would further increase the denominator (and not affect the numerator), leading to a lower MVPF associated with audits. Finally, expanding audits on the rich (poor) could lead to a decrease (increase) in income, which in principle would reduce (increase) the net revenue raised from audits and thus raise (lower) the MVPF. While our baseline results do not provide evidence for the idea that individuals who are audited choose to lower their incomes, we note that ex ante responses could differ from our within-person responses after being audited. The estimation of these additional components of the MVPF are important directions for future work.

E.5 Comparing the MVPF of Audits to the MVPF of Tax Rate Changes

As noted in Hendren and Sprung-Keyser [2020], for any two small policy changes, increasing spending on policy 1 financed by raising revenue from policy 2 will increase social welfare if and only if $\eta_1 MVPF_1 > \eta_2 MVPF_2$, where η_j is the social marginal utilities of income of the beneficiaries of policy j (i.e. giving \$1 to these beneficiaries raises social welfare by η_j). Therefore, we can evaluate the relative desirability of expanding audits versus increasing taxes as a method of raising revenue by comparing the MVPF of expanded audits to the MVPF of tax changes. To do so, it is natural to extend the model above to allow for heterogeneity in income choices and think about these MVPFs separately across the income distribution. Incorporating heterogeneity in incomes into the model is easily introduced by allowing the disutility of earnings to vary across individuals, which we index by θ, $\psi\left(y_t;\theta\right)$. The distribution of types in the population in turn generates an income distribution.

We can also compute the MVPF using the formula above conditional on income. If audit rates depend on income, individuals may change their incomes to reduce the probability of audit. For example, increasing the audit probability on top earners could cause people to reduce their reported taxable income, thus increasing the effective cost of the audit (by $T'\left(y_t\right)\frac{dy_t}{dp_t}$), and subsequently increasing the MVPF of the audit. This could be captured by including this term in the denominator of the MVPF. For our calculations, we do not make any such adjustment both because we do not have an empirical estimate of this potential behavioral response and also because in practice taking actions to reduce one's income in hopes of preventing an audit can increase the likelihood of the audit. In addition, in our event studies we find no evidence that audits cause reductions in future incomes.

As a result, we can write the MVPF of expanded audits around a given point of the income distribution as

$$MVPF^{Audit}(y) = \frac{R(y) + B(y)}{R(y) - C(y)}$$

where $R(y)$ is the average revenue per audit of those with incomes near y, $C(y)$ is the marginal cost of audits for those with incomes near y, and $B(y)$ is the average taxpayer burden of audits for those with incomes near y. The purple circles in Figure IX report these estimates of the MVPFs of tax audits by decile of the income distribution.

We can now compare the MVPF of tax audits to the MVPF of changes in the income tax schedule across the income distribution. In the environment above, Hendren (2020) shows that the MVPF of a tax change targeted to a particular region of the income distribution is given by

$$MVPF^{Tax}(y) = \frac{1}{1 + FE(y)}$$

where $FE(y)$ is the impact of the behavioral response to a small tax cut targeted to those earning near y on the government budget. Under quasilinear utility, this is given by

$$FE(y) = \frac{T'(y)}{1 - T'(y)} \kappa(y) \epsilon^c(y)$$

where $\epsilon^c(y)$ is the compensated elasticity of taxable income with respect to the marginal tax rate, $T'(y)$ is the marginal tax rate at y, and $\kappa(y)$ is the local Pareto parameter of the income distribution. The triangles in Figure IX shows the shape of $MVPF^{Tax}(y)$ as constructed in Hendren [2020].

E.6 Heterogeneous Welfare Weights and Comparisons Across Policies

Our key finding is that the MVPF for audits of top earners is lower than the MVPF for of tax cuts for income earners, which suggests expanded audits would raise revenue at lower welfare cost than tax rate increases. Formally, raising revenue from expanded audits and using it to finance lower taxes raises welfare if and only if

$$\eta^{audits} MVPF^{audits} < \eta^{tax} MVPF^{tax}$$

where η^{audits} is the incidence-weighted average social welfare weight for those who are audited and η^{tax} is the incidence-weighted average welfare weight of those facing the tax change. Formally, if R_i is the revenue raised from individual i and B_i is the "burden" of the audit for individual i, then $R_i + B_i$ is individual i's willingness to pay for the audit. So, if η_i is individual i's social welfare weight (i.e. social welfare goes up by η_i if we give \$1 to individual i), we have

$$\eta^{audits} = \frac{\sum_i (R_i + B_i) \eta_i}{\sum_i (R_i + B_i)}$$

For example, it could be natural to imagine that the social planner places different weight on compliant and noncompliant taxpayers. Let η_d denote the welfare weight for noncompliant taxpayers and η_h the welfare weight for compliant taxpayers. Let d denote the fraction of noncompliant taxpayers amongst the set of marginally audited individuals. We can then write

$$\eta^{audits} = \frac{d\left(B_d + \frac{R}{d}\right)\eta_d + (1-d)\, B_h \eta_h}{B+R}$$

where R is the revenue per audit, so that R/d is the revenue per audit of a noncompliant taxpayer (assume no revenue from compliant taxpayers), and $B = dB_d + (1-d)\, B_h$ is the average taxpayer burden among those audited.

If the social planner places little weight on noncompliant taxpayers, $\eta_d \approx 0$, this will reduce the welfare weight of the audited group to the extent that it contains noncompliant taxpayers (e.g. $d = 1$). For this reason, we reason in the main text that $\eta^{audits} < \eta^{tax}$ so that the fact that $MVPF^{audits} < MVPF^{tax}$ suggests it is more efficient to raise revenue from audits than taxes at the top of the income distribution. A potential countervailing force that is not immediately captured in our model is the fact that auditing compliant taxpayers could impose significant psychological and financial burdens on those taxpayers. In our baseline approach, we quantify the taxpayer burden using the average reported time cost of complying with an audit, but with additional psychological costs B_h could be high.

We can use our results to ask how high B_h would need to be to overturn the conclusion that expanding audits raise revenue at lower welfare cost than raising tax rates. To do so, recall that expanding audits raises revenue at lower welfare cost than raising tax rates if and only if $\eta^{audits} MVPF^{audits} < \eta^{tax} MVPF^{tax}$. The equation above provides η^{audits}. For tax rate increases the welfare weight is just the population welfare weight at the relevant income level (e.g. average welfare weight of top earners). This is the average across compliant and noncompliant taxpayers. Let p denote the fraction of noncompliant taxpayers in the total population (note $p < d$ when audits are targeted more towards noncompliant taxpayers). Then, $\eta^{tax} = p\eta_d + (1-p)\,\eta_h$. So, we have that $\eta^{audits} MVPF^{audits} < \eta^{tax} MVPF^{tax}$ if and only if:

$$\frac{\frac{(1-d)B^h\eta_h + \left(R + dB^d\right)\eta_d}{R+B}}{\eta^h(1-p) + \eta^d p} < \frac{MVPF^{tax}}{MVPF^{audits}} = \frac{MVPF^{tax}}{(R+B)\,/\,(R-C)}$$

where B is the average per audit burden and R is the average per-audit revenue.

If the social planner puts little welfare weight on noncompliant taxpayers, $\eta^d \approx 0$, then audit expansions are preferred to tax rate increases if and only if

$$\frac{\frac{(1-d)B^h}{R+B}\eta_h}{\eta^h(1-p)} < \frac{MVPF^{tax}}{(R+B)\,/\,(R-C)}$$

$R + B$ cancels on both sides so that we are left with

$$\frac{(1-d)B^h}{(1-p)} < MVPF^{tax}\,(R-C)$$

or

$$B^h < MVPF^{tax}\,(R-C)\,\frac{1-p}{1-d}$$

In other words, the planner wishes to increase audits as long as B^h is less than the MVPF of the tax multiplied by the net revenue of the tax multiplied by the ratio of the fraction of compliant taxpayers overall relative to compliant taxpayers among those audited. Taking estimates from the literature that find the MVPF of increasing top tax rates is around 1.5, assuming most taxpayers are compliant, e.g. $p \approx 0.9$, letting d be the fraction of audited taxpayers that see a change in their tax liability (which is around 60%), and using our estimate of the net revenue raised from an audit of a taxpayer in the top 0.1% of the income distribution of around \$103,000 yields

$$B^h < 1.5 * \$103,000 * \frac{0.9}{0.4} = \$348,000$$

In other words, unless those who are audited and found to owe no additional tax are willing to pay more than \$300,000 to avoid the audit, the social planner would wish to raise revenue by expanding audits rather than by raising tax rates on top earners.

E.7 Audit Rates versus Audit Penalties

A large literature on considers the optimal mix of audit rates and audit penalties (Becker, 1968; Allingham et al., 1972), and how that mix varies across the income distribution (Border and Sobel, 1987). While optimal audit penalties are beyond the scope of our empirical analysis, here we illustrate how one can use the MVPF framework to think about the welfare impact of changes in audit penalties relative to increases in audit probabilities. In particular, we first show how one can replicate the classic Becker (1963) and Allingham and Sandmo (1972) result that it is optimal to have high penalties and low audit probabilities when audits are costly. We then discuss briefly how the MVPF provides an empirical path forward for generalizing this conclusion to settings where assumptions made in those models may not hold.

To gain intuition for the MVPF of changes in audit penalties, we can return to the model above and suppose that the government is considering modifying the penalties people must pay in some period. To be specific, suppose the government is considering raising the penalties paid in period 1 to add an additional penalty ν per dollar of money that is evaded. In other words, suppose penalties in period 1 are now given by $\phi^{\alpha_1}(e) = \phi^{\alpha_1}(e) + \nu e$, where e is the amount evaded and ν parameterizes an increase in the evasion penalty rate. With this modification, the FOC for evasion now becomes

$$\frac{\partial \phi^{\alpha_t}}{\partial e}\left(e^*_{t,\alpha_t}\right) = \frac{1-p_t}{p_t} - \nu$$

Assuming penalties are convex, then the increase in ν will reduce evasion, $de^*_{t,\alpha_t}/d\nu < 0$. The question we ask is: what is the welfare impact of changes in penalties relative to changes audit probabilities.

To assess this, note that the social welfare impact of increasing audit penalties can be written as $\eta^{penalty} MVPF^{penalty}$, where $\eta^{penalty}$ is the average social marginal utility of income of those who are subjected to the higher penalties and $MVPF^{penalty}$ is the MVPF of higher penalties, given by the ratio of people's willingness to pay to avoid the penalties relative to the net revenue raised by those penalties (which includes both changes in penalty revenue and also changes in voluntary tax revenue paid from deterrence effects). Raising revenue by raising penalties instead of raising audit rates will be preferred whenever

$$\eta^{audits} MVPF^{audits} > \eta^{penalty} MVPF^{penalty} \tag{9}$$

In Allingham and Sandmo (1972) and Becker (1968), there is a representative agent so that there is no difference in social welfare weights, $\eta^{audits} = \eta^{penalty}$. The key question is then: what is the MVPF of raising penalties? To start, note that an increase in penalties in period 1, $d\nu$, will generate a willingness to pay to avoid those increased penalties of $e^*_{1,\alpha_1} d\nu$. The revenue generated from the increased penalties in period 1 will be the sum of the mechanical revenue generated, $e^*_{1,\alpha_1} d\nu$, plus the impact of the behavioral response to the penalties on the government budget (i.e. the general deterrence effect). However, because penalties are entirely financial, individuals' evasion choices minimize government revenue. Hence, the behavioral response to the change in ν has no effect on government revenue. As a result, the cost to the government is simply the mechanical revenue, meaning the MVPF is given by

$$\begin{aligned} MVPF^{penalty} &= \frac{e^*_{1,\alpha_1} d\nu}{e^*_{1,\alpha_1} d\nu + \frac{de^*_{1,\alpha_1}}{d\nu}\left[(1-p_1) - p_1 \frac{\partial \phi^{\alpha_1}}{\partial e}\left(e^*_{1,\alpha_1}\right)\right]} \\ &= \frac{e^*_{1,\alpha_1} d\nu}{e^*_{1,\alpha_1} d\nu} \\ &= 1 \end{aligned}$$

In contrast, recall that the MVPF of expanded audit rates is always greater than 1: individuals are willing to pay the mechanical revenue to avoid the audit, but the revenue to the government is lower due to the net cost of expanded audits. Hence, $MVPF^{penalty} = 1 < MVPF^{audits}$. This is the classic logic of Becker (1968): raising penalties raises revenue more efficiently than raising audit rates.

This conclusion rests on strong assumptions about how taxpayers respond to changes in audit rates and penalties and about the social planner's preferences. As noted above in Section E.4.2, the MVPF framework can incorporate cases where general deterrence has first order impacts on government revenue – given empirical estimates of the deterrence effects of raising audit rates and raising penalties, one can construct more general MVPF measures of these policies that allow for general deterrence to affect government revenue. In addition, the MVPF framework is not limited to a representative agent and can place differential welfare weights on audited individuals — social planners may find it undesirable to give very low consumption to those being fined. Given estimates of the willingness to pay to avoid the higher penalties combined with the net revenue raised by the penalties, one can compare the relative desirability of expanded audits versus higher penalties using Equation 9. In this sense, the MVPF framework provides a path for replacing the theoretical assumptions with empirical elasticities on how audit design affects behavior and revenue.

References

Allingham, M. G., A. Sandmo, et al. (1972). Income tax evasion: A theoretical analysis. *Taxation: critical perspectives on the world economy 3*(1), 323–338.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of political economy 76*(2), 169–217.

Border, K. C. and J. Sobel (1987). Samurai accountant: A theory of auditing and plunder. *The Review of economic studies 54*(4), 525–540.

Hendren, N. (2020). Measuring economic efficiency using inverse-optimum weights. *Journal of public Economics 187*.

Hendren, N. and B. Sprung-Keyser (2020). A unified welfare analysis of government policies. *The Quarterly Journal of Economics 135*(3), 1209–1318.

Holtzblatt, J. and J. McGuire (2020). Effects of recent reductions in the internal revenue service's appropriations on returns on investment. In *IRS Research Bulletin: Proceedings of the 2020 IRS/TPC Research Conference*.

Keen, M. and J. Slemrod (2017). Optimal tax administration. *Journal of Public Economics 152*, 133–142.

Sarin, N. and L. H. Summers (2019). Shrinking the tax gap: approaches and revenue potential. Technical report, National Bureau of Economic Research.