Tax Evasion at the Top of the Income Distribution: Theory and Evidence *

John Guyton
Internal Revenue Service

Patrick Langetieg
Internal Revenue Service

Daniel Reck
London School of Economics

Max Risch
Carnegie Mellon University

Gabriel Zucman
University of California Berkeley

February 19, 2021

First Draft. Comments Welcome.

Abstract

This paper studies tax evasion at the top of the U.S. income distribution using IRS micro-data from (i) random audits, (ii) targeted enforcement activities, and (iii) operational audits. Drawing on this unique combination of data, we demonstrate empirically that random audits underestimate tax evasion at the top of the income distribution. Specifically, random audits do not capture most tax evasion through offshore accounts and pass-through businesses, both of which are quantitatively important at the top. We provide a theoretical explanation for this phenomenon, and we construct new estimates of the size and distribution of tax noncompliance in the United States. In our model, individuals can adopt a technology that would better conceal evasion at some fixed cost. Risk preferences and relatively high audit rates at the top drive the adoption of such sophisticated evasion technologies by high-income individuals. Consequently, random audits, which do not detect most sophisticated evasion, underestimate top tax evasion. After correcting for this bias, we find that unreported income (as a fraction of true income) rises from 7% in the bottom 50% to more than 20% in the top 1%. We estimate that 42% of federal income taxes unpaid are owed by the top 1%; collecting all unpaid federal income tax from this group would increase federal revenues by about $200 billion annually.

*Corresponding Author: Daniel Reck, d.h.reck@lse.ac.uk. *We thank Gerald Auten, Brian Galle, Bhanu Gupta, Tom Hertz, Xavier Jaravel, Drew Johns, Barry Johnson, Camille Landais, Katie Lim, Emily Lin, Larry May, Alicia Miller, Annette Portz, Peter Rose, Emmanuel Saez, Brenda Schafer, Clifford Scherwinski, Joel Slemrod, Matt Smith, Johannes Spinnewijn, David Splinter, Alex Turk, and Alex Yuskavage for helpful discussion, support, and comments on preliminary versions of this work. Jeanne Bomare and Baptiste Roux provided excellent research assistance. All remaining errors are our own. Financial support from the Washington Center for Equitable Growth, the Stone foundation, Arnold Ventures, and the Economic and Social Research Council is gratefully acknowledged.*

The views expressed here are those of the authors and do not necessarily reflect the official view of the Internal Revenue Service. This project was conducted through the Joint Statistical Research Program of the Statistics of Income Division of the IRS. All data work for this project involving confidential taxpayer information was done at IRS facilities, on IRS computers, by IRS employees, and at no time was confidential taxpayer data ever outside of the IRS computing environment. Reck and Risch are IRS employees under an agreement made possible by the Intragovernmental Personnel Act of 1970 (5 U.S.C. 3371-3376).

1 Introduction

How much do wealthy individuals evade in taxes? And what are the main forms of tax noncompliance of the top of the income distribution? Because taxable income and tax liabilities are highly concentrated at the top of the income distribution, understanding noncompliance by high-income taxpayers is critical for the analysis of tax evasion, for tax enforcement, and for the conduct of tax policy.

A key difficulty in studying tax evasion by the wealthy is the complexity of the forms of tax evasion at the top, which can involve legal and financial intermediaries, sometimes in countries with a great deal of secrecy. This complexity means that one single data source is unlikely to uncover all forms of noncompliance at the top. In this paper, we attempt to overcome this limitation in the U.S. context by combining a wide array of sources of micro data, including (i) random audit data, (ii) the universe of operational audits conducted by the IRS, and (iii) targeted enforcement activities (e.g., on offshore bank accounts). Drawing on this unique combination of data, we show that random audits underestimate tax evasion at the top-end of the income distribution. We provide a theoretical explanation for this fact, and we propose a methodology to improve the estimation of the size and distribution of tax noncompliance in the United States.

The starting point of our analysis is the IRS random audit program, known as the National Research Program. Random audits are the most common tool used to study and measure the extent of tax evasion. Researchers use random audits to test theories of tax evasion (Kleven et al., 2011), and tax authorities use them to estimate the overall extent of tax evasion and target audits strategically (IRS, 2019). The idea of the random audit as the gold standard for understanding tax evasion comes from the traditional appeal of random sampling, combined with the classic deterrence model of tax evasion (Allingham and Sandmo, 1972), an implicit assumption of which is that audits lead to the detection of all tax evasion. In the real world, however, random audits do not detect all forms of tax evasion. Random audits are well designed to detect common forms of tax evasion, such as unreported self-employment income, overstated deductions, and the abuse of tax credits. But, we argue, these audits may not detect sophisticated evasion strategies, because doing so can require much more information, resources and specialized staff than available to tax authorities for their random audit programs.

Our first contribution is to document and quantify the limits of random audits when it comes to detecting top-end evasion in the United States. We find that detected evasion declines sharply at the very top of the income distribution, with only a trivial amount of evasion detected in the top 0.01%. Our analysis uncovers two key limitations of random audits which can account for much of this drop-off: tax evasion through foreign intermediaries (e.g., undeclared foreign bank accounts) and tax evasion via pass-through businesses (e.g., partnerships). First, we find that offshore tax evasion goes almost entirely undetected in random

audits.[1] To establish this result, we analyze the sample of U.S. taxpayers who disclosed hidden offshore assets in the context of specific enforcement initiatives conducted in 2009–2012. A number of these taxpayers had been randomly audited just before this crackdown on offshore evasion. In over 90% of these audits, the audit had not uncovered any foreign asset reporting requirement, despite the fact that these taxpayers did own foreign assets. Second, we find that tax evasion occurring in pass-through businesses (partnerships and S-corporations, businesses whose ownership is highly concentrated) is substantially under-detected in random audits. Examiners are often limited in their ability to check the degree to which pass-through businesses have duly reported their income, especially for the most complex businesses. Thus, while the income of taxpayers in the bottom 99% of the income distribution is comprehensively examined, up to 35% of the income earned at the top goes un-examined in the context of random audits.

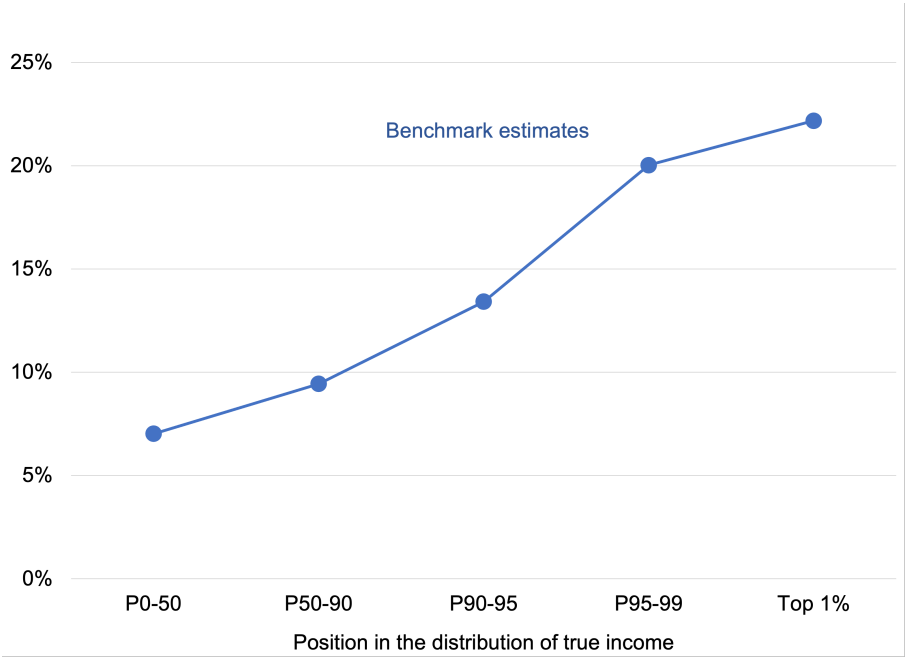
Our second contribution is to propose improved estimates of how much income (relative to true income) the various groups of the population under-report—and to investigate the consequences of this under-reporting for the measurement of inequality. We do so by starting from evasion estimated in random audits and proposing a correction for sophisticated evasion that goes undetected in these audits. Figure 1 depicts our benchmark results. We find that total under-reported income rises from about 7% of true income in the bottom 50% of the income distribution to 22% in the top 1%. Out of this 22%, close to 6 percentage points correspond to sophisticated evasion that goes almost undetected in random audits. We show that accounting for under-reported income increases the top 1% fiscal income share in the United States significantly. Using the same methodology as Johns and Slemrod (2010) (which does not include the correction for sophisticated evasion we propose in this paper) and applying it to random audit data from more recent years, we find that accounting for tax evasion increases the top 1% fiscal income share from 20.3% before audit to 22.2% on average over 2006–2013. When factoring in our correction for sophisticated evasion, the top 1% fiscal income share further increases to 22.8%. The result that accounting for tax evasion increases inequality is robust to a wide range of robustness tests and sensitivity analysis (for instance, it is robust to assuming zero offshore tax evasion).

Our third contribution is to explain why general-purpose random audits are not uniformly able to detect noncompliance across the income distribution. We present a model in which high-income taxpayers adopt sophisticated evasion strategies. We show that introducing this element in the canonical Allingham and Sandmo (1972) tax evasion model changes our understanding of tax evasion by high-income persons.

The model allows a taxpayer to adopt some costly form of tax evasion that is unlikely to be discovered on audit at some cost. We show that adoption of such an evasion technology is likely to be concentrated

[1] Our data cover the period prior to the collection of third-party reported information on foreign bank accounts, which started in 2014; we analyze how our results can inform knowledge about post-2014 evasion in Section 4.

FIGURE 1: UNREPORTED INCOME (AS A PERCENT OF TRUE INCOME) BY INCOME GROUP



Note: This figure shows unreported income as a fraction of true income across the distribution of true income in the United States in 2006–2013. P0-50 denotes the bottom 50% tax units, P50-90 the next 40%, etc. Income is market income (defined as total income reported on form 1040 minus Social Security benefits, unemployment insurance benefits, alimony, and state refunds); tax units are ranked by their corrected (i.e., reported plus unreported) market income.

at the top of the income distribution for two reasons. First, high-income taxpayers have a greater demand for sophisticated evasion strategies that reduce the probability of detection if (i) the desired rate of evasion does not become trivial at large incomes, and (ii) the cost of adopting becomes a trivial share of income at large incomes. This is true even holding the probability of audit by income fixed. Second, overall audit rates and scrutiny of tax returns are substantially higher at the top than at the bottom of the distribution, making evasion that is less likely to be detected and corrected on audit more attractive at the top. We can also re-interpret the model to think about situations where the outcome of an audit, if it occurs, is uncertain. With this interpretation, for the same reasons as before, we show that high-income people are then more likely to adopt positions in the "gray area" between legal avoidance and illegal evasion. From the point of view of the tax authority, we show theoretically that high resource costs of pursuing sophisticated forms of tax evasion, such as protracted litigation or more sophisticated audits of a complex network of closely-held businesses, can pose practical limits on the extent to which the tax authority can pursue these types of tax evasion by high-income people. This is especially the case when resource constraints are exogenous and not changed when sophisticated evasion becomes more prevalent.

These findings have implications for the academic literature, for policymakers, and for the public debate

over income taxes at the top. Academically, our findings show that the existing framework for thinking about tax evasion has limitations when it comes to top-end tax evasion. The increasingly conventional wisdom is that taxpayers seldom evade taxes supported by third-party information (Kleven et al., 2011; Carrillo et al., 2017; Slemrod et al., 2017; IRS, 2019), and that deterring evasion where taxes are not supported by third-party information requires increasing the audit rate, or the penalty rate, or, arguably, increasing tax morale (Luttmer and Singhal, 2014). This characterization works well for the middle and bottom of the income distribution. However, it misses the importance of deliberate concealment of income (even from auditors) at the top, and the adoption of aggressive interpretations of tax law for sheltering purposes. From a government revenue perspective, the top of the income distribution is the sub-population where understanding the extent of tax evasion is the most important, due to the high and increasing concentration of income in the United States (Piketty and Saez, 2003; Piketty et al., 2018).

From a policy perspective, our results highlight that there is substantial evasion at the top which requires administrative resources to detect and deter. We estimate that 42% of federal income taxes unpaid are owed by the top 1% and that collecting all unpaid federal income tax from this group would increase federal revenues by about $200 billion annually. There has been much discussion in the United States about the fact that the audit rate at the top of the income distribution has declined. Our results suggest that such low audit rates are not optimal. As standard audit procedures can be limited in their ability to detect some forms of evasion by high-income taxpayers, additional tools should be mobilized to effectively combat high-income tax evasion. These tools include facilitating whistle-blowing that can uncover sophisticated evasion (which helped the United States start to make progress on detection of offshore wealth) and specialized audit strategies like those pursued by the IRS's Global High Wealth program and other specialized enforcement programs.[2] The IRS already uses these strategies, but resource constraints limit its capacity to do so (see, e.g., TIGTA, 2015). Our results suggest that investing in improved tools and increasing resources to support tax administration at the top of the distribution could generate substantial tax revenue (a point also made by, e.g., Sarin and Summers, 2020). Additionally, our results suggest that incorporating insight from sources of data beyond conventional random audits may be useful for risk assessment, audit selection, and the allocation of IRS resources to alternative types of enforcement.

The rest of this paper is organized as follows. Section 2 studies the distribution of noncompliance in random audit data. Section 3 provides direct evidence that some forms of evasion are (i) highly concentrated at the top of the income distribution, (ii) effectively invisible in random audits, and (iii) quantitatively important for the measurement of income at the top. In Section 4 we present our new estimates of the distribution of noncompliance and we investigate their implications for the measurement of inequality. Section 5

[2] See https://www.irs.gov/irm/part4/irm_04-052-001 for information on the Global High Wealth program.

presents our theory of why some noncompliance goes undetected, and Section 6 concludes.

2 The Distribution of Noncompliance in Random Audits

The National Research Program (NRP) random audits are the main data source used to study the extent and nature of tax evasion in the United States (see, e.g., Andreoni et al., 1998; Johns and Slemrod, 2010; IRS, 2016, 2019; DeBacker et al., 2020).[3] NRP auditors assess compliance across the entire individual tax return—the Form 1040—based on information from the schedules of the Form 1040, third-party information reports, the taxpayer's own records, and measures of risk comparing all this information to information on the broader filing population.[4] The most commonly cited statistics from random audit studies are estimates of the *income under-reporting gap*—the amount of income under-reported, expressed as a fraction of true income—and of the *tax gap*—the amount of tax that is legally owed but not paid, expressed as a fraction of the amount of tax legally owed. It has long been acknowledged that in the context of a random audit, some noncompliance may go undetected. The IRS uses a methodology, known as detection-controlled estimation (DCE), to estimate undetected noncompliance. Official IRS estimates (presented in, e.g., IRS, 2016, 2019) of the aggregate tax gap use the DCE methodology, as do existing estimates of the distribution of unreported income and evaded income taxes (Johns and Slemrod, 2010).

In this section, we start by describing evasion seen in raw NRP random audits without any correction for undetected noncompliance (in particular, before DCE correction), and then show results including the DCE correction.[5] All our analyses pool data from the NRP random audits conducted in tax years 2006–2013. The NRP uses a stratified random sample which over-samples top earners to ensure good coverage at the top. The pooled sample we use includes 105,167 audited taxpayers, of which 12,003 in the top 1% of the reported income distribution. We use the NRP weights to compute statistics that are representative of the full population. Our sample is large enough to obtain precise estimates for groups as small as the top 0.01% (although splitting this very top group by other characteristics tends to leave us with too little statistical power for informative analysis).

[3] Further background on the NRP is in the Internal Revenue Manuals here: https://www.irs.gov/irm/part4/irm_04-022-001. We use the term evasion in this paper to refer to unintentional and intentional noncompliance with tax obligations. We do not attempt to distinguish between intentional evasion and unintentional noncompliance and acknowledge that the boundary between these is fuzzy.

[4] Earlier IRS random audit studies under the Taxpayer Compliance Measurement Program (TCMP) consisted of line-by-line examinations of the individual tax return. The NRP aims to provide a similarly comprehensive measure of compliance at a reduced administrative cost and burden on the taxpayer. See Brown and Mazur (2003) for more on the TCMP and how the NRP uses revised procedures to achieve similar objectives.

[5] Recent work by DeBacker et al. (2020) performs similar analysis to our work here (in Section 2.1), without DCE. Our results in this section are similar to theirs, but because of our subsequent results, our interpretation of these patterns is very different. Specifically, we argue that the low detected evasion at the top is a consequence of the fact that evasion at the top is less likely to be detected in a random audit, not that high-income taxpayers are much more compliant.

2.1 The Distribution of Detected Evasion

To begin with, we take the NRP random audit data at face value (i.e., before DCE correction) and estimate income under-reported as a fraction of audit-corrected income and tax evaded as a fraction of tax due, within income groups defined based on audit-corrected income.[6]

On aggregate, NRP audits find that 3.9% of true income is under-reported and 7.9% of taxes owed are not paid, before any correction for undetected evasion.[7] Note that these numbers are significantly smaller than the official IRS (2016, 2019) estimates of noncompliance, in which 13.9% of aggregate true income is found unreported and 20.3% of individual income taxes owed are found unpaid (see Section 2.2), because the official estimates factor in the DCE adjustment described below.

Figure 2a shows that the detected income under-reporting gap hovers around 4% to 5% through most of the income distribution and falls sharply within the top 0.1% to less than 1% in the top 0.01%. Figure 2b decomposes unreported income by type of income, and Appendix Table A1 reports the composition of under-reported income in the full population and in the top 1%. We observe that by far the largest component of detected under-reporting involves corrections to Sole Proprietor income, on Schedule C of the individual income tax return. Under-reporting in this category comprises 50% of all evasion detected in the NRP. The next-largest category involves corrections to Form 1040 Line 21 income ("Other income"), which mostly reflect disallowed net-operating loss carryforwards and carrybacks. Adjustments to other income on line 21 are relatively uncommon but they can be large when they occur. The individuals with the largest adjustments in this category usually had negative reported income due to business losses in other years, but once those losses are disallowed, their audit-corrected income falls in the top 1% of the income distribution.

Appendix Figure A1 shows tax evaded as a fraction of tax due. The patterns are consistent with those seen in Figure 2a. At the very top, the detected tax gap is negligible. At the bottom of the distribution, taxes evaded appear large as a fraction of taxes owed. This is in part a mechanical consequence of the fact that refundable tax credits reduce taxes owed to close to zero (sometimes less than zero) for low-income taxpayers: by definition, groups with close to zero "net" taxes owed are bound to have very high under-reporting rates relative to taxes owed.[8]

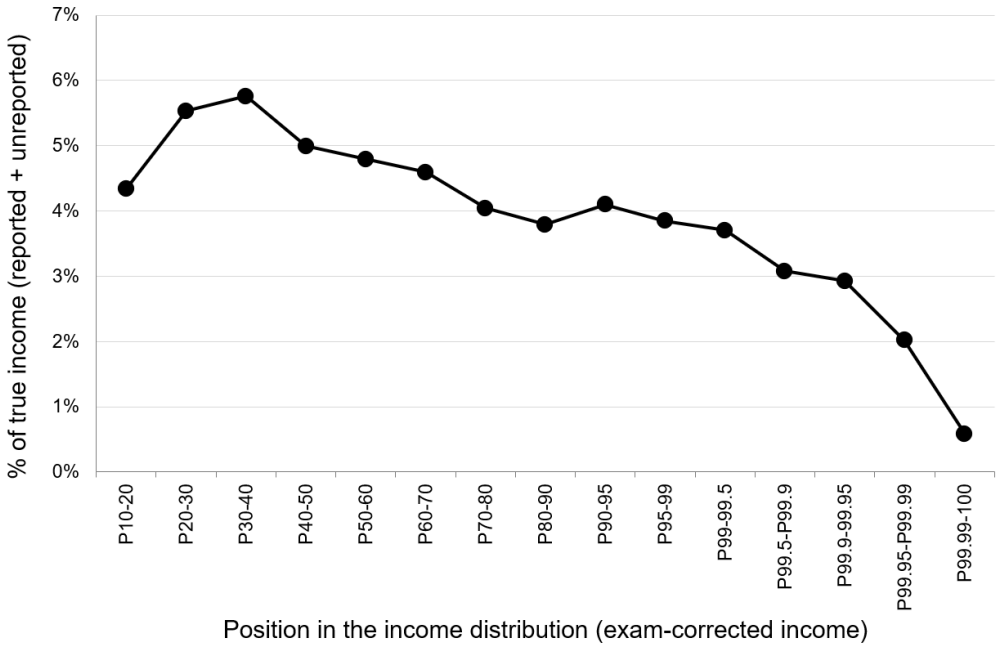
[6] Unless otherwise noted, all our analyses in this paper rank taxpayers by their estimated true income (with different measures of "true income" depending on the method used to estimate unreported income). Ranking taxpayers by reported income leads to downward-biased estimates of the income under-reporting gap and tax gap at the top, because taxpayers with high reported income are selected on high compliance (they declare high incomes).

[7] These numbers are for the entire population and include taxpayers who were found over-reporting income. Taxpayers who under-reported income under-reported 4.2% of aggregate true income, while taxpayers who over-reported income over-reported the equivalent of 0.3% of aggregate true income. The majority of over-reported income is in the bottom half of the exam-corrected income distribution; the implications of over-reporting for aggregate tax liabilities and for noncompliance at the top are negligible. Unless otherwise noted, our computations in this paper include all taxpayers, including those found over-reporting income.

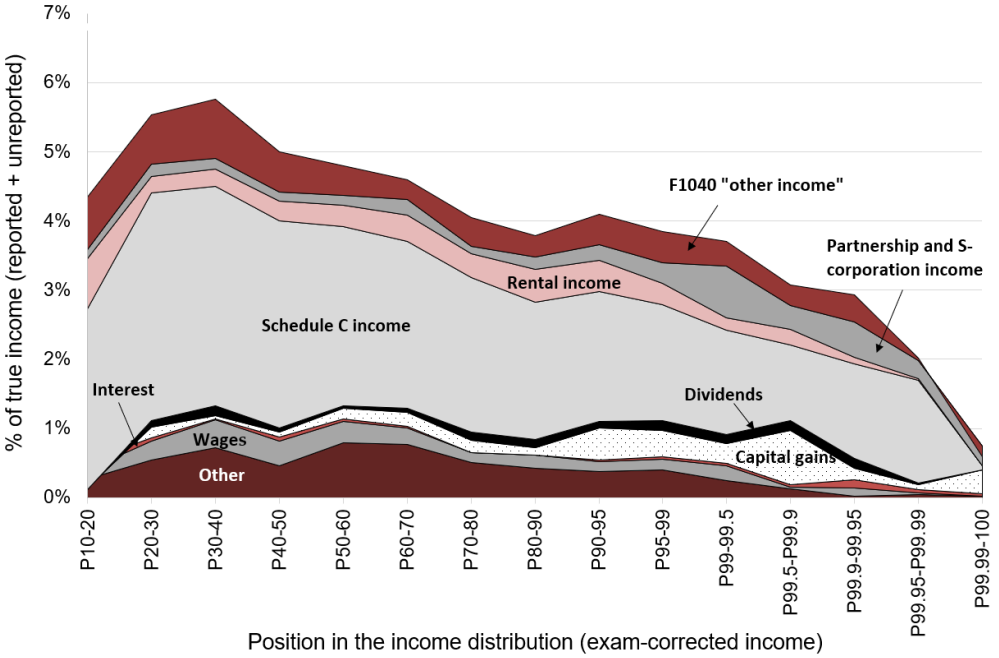
[8] A better way to study the tax gap is to consider that taxes owed must be non-negative (to avoid dividing taxes evaded by zero), i.e., to treat the refundable portion of tax credits as government transfers (instead of negative taxes). Figure A2 shows that excluding tax credits from both reported and corrected taxes reduces the tax gap at the bottom by about 7 percentage points. A fuller analysis would include all taxes (or at least all federal taxes) in the analysis, including most importantly payroll taxes. Because of the relatively

FIGURE 2: UNREPORTED INCOME DETECTED IN RANDOM AUDIT DATA BEFORE DCE CORRECTION

(a) Unreported Income (% of True Income)



(b) Decomposition by Type of Income



Notes: This Figure shows the pattern of income under-reporting uncovered in NRP random audit data for 2006-2013, without any correction for undetected evasion (in particular before DCE correction). Tax units are ranked by their exam-corrected market income (defined as total income reported on form 1040 minus Social Security benefits, unemployment insurance benefits, alimony, and state refunds). We observe that detected unreported income decreases sharply within the top 1% of the income distribution. Mis-reporting of Schedule C income comprises the bulk of evasion detected in NRP random audits. We also observe that by contrast, very little evasion is detected for partnership and S-corporation business income, which is an important source of income at the top.

Three remarks are in order. First, the NRP detects little evasion on wages, interest, dividends, and capital gains. A natural explanation for this fact is that these forms of income are subject to substantial third-party reporting (IRS, 2016; Kleven et al., 2011).[9] However, interest, dividends, and capital gains accruing to offshore accounts only started being subject to information reporting with the implementation of the Foreign Account Tax Compliance Act (FATCA) in 2014, after our period of study. Thus, the low evasion rates on financial capital income recorded in the 2006–2013 NRP may, in part, be due to the fact that some evasion on offshore capital income went undetected.

Second, there is an asymmetry in the estimated rates of detected evasion across different types of business income. The estimated detected income under-reporting rate for sole proprietor income, which is supported by relatively little third-party information, is about 37% overall and 24% for the top 1%.[10] By contrast, the estimated detected under-reporting rate for pass-through business income (i.e., partnerships and S-corporations business income) is 4.2% overall and just 2.0% for the top 1%. Pass-through business income is common at the top of the income distribution, and like sole proprietorship income, is supported by relatively little third-party information. As we shall see, sole proprietor income is subject to extensive examination in the context of the NRP, while the examination of pass-through business income faces practical limits described below.

Third, we estimate that detected evasion among the very wealthy is extremely low in the NRP. In the top group we consider, the top 0.01% (by exam-corrected income), we estimate that just 0.6% of true income is under-reported. This is a direct consequence of the fact that the NRP uncovers little noncompliance on interest, dividends, pass-through business income, and capital gains—the key sources of income in the top 0.01%.

2.2 The Distribution of Evasion with the DCE Methodology

The IRS tax gap estimates acknowledge that some evasion may go undetected in the context of NRP random audits by employing a technique called Detection Controlled Estimation (DCE), under which detected evasion is scaled up to account for undetected evasion. DCE methodology is based on Feinstein (1991). The detection process is modeled by positing that, conditional on evasion occurring, only a fraction is detected depending on the characteristics of the return examined (presence of self-employment income, schedules filed, etc.) and of the examiner (experience, age, etc.). Feinstein (1991) estimates such a model by maximum

low rate of evasion on payroll taxes (IRS, 2019), this would reduce the ratios of taxes evaded to taxes owed shown in Figure A1 in the bottom and middle of the distribution significantly. We leave this task to future research.

[9] Capital gains were subject to some information reporting throughout our period of study; a reform added the requirement that brokers report not just sale price but also cost basis starting in tax year 2011.

[10] See Slemrod et al. (2017) for more information on the limited third-party reporting for sole proprietor income. In official Tax Gap statistics the estimate under-reporting rate for sole proprietor income is 56% (IRS, 2019). The difference is attributable to the DCE adjustment.

likelihood and finds that about a third of tax evasion goes detected (i.e., if all examiners were as perceptive as the examiners who uncover the most evasion, three times more evasion would be detected). To adjust for unreported income that examiners were unable to detect, the IRS applies DCE to the returns subject to audit. Separate multipliers are applied for low-visibility and high-visibility income and for taxpayers with reported total positive income above and below $100,000.[11] The same approach is followed by Johns and Slemrod (2010) to study the distribution of noncompliance in 2001.

Figure 3 shows the distribution of noncompliance after applying the DCE methodology in 2006–2013. A number of results are worth noting. First, the DCE adjustment roughly triples estimated detected noncompliance. After DCE correction, we estimate that 13.9% of true income was under-reported on average over the 2006–2013 period, and 20.3% of individual income taxes owed were not paid. These numbers are essentially identical to the estimates of the tax gap reported in IRS (2019, Figure 1) for the years 2011-2013.[12] Second, the DCE adjustment for the most part reverses the decreasing profile of raw detected tax evasion by income group seen in Figure 2. Income under-reported now rises from about 7% of DCE-corrected income at the bottom of the distribution to about 25% in the lower half of the top 1%. A key reason for this rising profile is re-ranking when audited taxpayers are ranked by their DCE-corrected income, as shown by Appendix Figure A3. For example, the DCE adjustment barely increases the very small amount of evasion detected among taxpayers in the top 0.01% by exam-corrected income. Because DCE adjustment increases the estimated amount of under-reported income for taxpayers initially below the top 0.01%, some of these taxpayers move to the top 0.01%. Third, and despite this re-ranking, even with DCE a steep drop-off in estimated evasion remains at the very top. Estimated under-reported income falls from just over 25% of true income between the 99^{th} and 99.5^{th} percentile to just over 20% in the top 0.5%.

The top panel of Figure 3 also compares our estimates of the distribution of noncompliance to those of Johns and Slemrod (2010), which factor in the same DCE methodology as the one we use in this paper. The difference between our work and Johns and Slemrod (2010) is that we use more recent waves of the NRP.[13] We also have larger sample sizes: 105,167 audited taxpayers in our sample, as opposed to 36,699 in the 2001 NRP used by Johns and Slemrod (2010). Overall, we find a broadly similar distribution of under-reported income and taxes evaded. A number of differences are worth noting. Under-reported income is higher on

[11] Total positive income is the sum of all positive amounts of the various components of income reported on an individual tax return, and thus excludes losses. Johns and Slemrod (2010) provide more details on DCE methodology as used in the 2001 wave of the NRP. DCE methods have been slightly revised in more recent tax gap studies (IRS, 2019), although the basic approach remains the same.

[12] The gross tax gap for tax filers is estimated in IRS (2019) at $283 billion per year, which is 20.2% of the estimated true individual income tax liability of $1,398 billion. An additional $31 billion gross tax gap is estimated for non-filers, based not on NRP data but on the Administrative Data Method; see IRS (2019, p. 15).

[13] As already noted, the DCE methodology has been slightly refined between the 2001 NRP used by Johns and Slemrod (2010) and the 2006–2013 waves of the NRP we use in this paper. However to maximize comparability, in this paper we perform the same DCE procedure as Johns and Slemrod (2010). The updated DCE methods, described in IRS (2019), are similar to the methods we employ; both are based on Feinstein (1991).

aggregate in our series (13.9% of true income) than in the 2001 NRP (11%). It is also more concentrated, with 30% of under-reported income going to the top 1% ranked by true income (see Table 2 below) as opposed to 27% in Johns and Slemrod (2010, Table 3). Third, accounting for under-reported income increases the top 1% income share in the 2006–2013 NRP, from 20.3% before audit to 22.2% (see Table 3 below). By contrast, under-reported income is distributionally neutral in Johns and Slemrod (2010, Table 5): adding it leaves the top 1% income share stable at 17.8% in 2001. Last, at the bottom of the distribution we find lower ratios of taxes evaded to taxes owed than Johns and Slemrod (2010).[14]

2.3 Limits of the DCE Methodology: A Direct Test

Although Detection Controlled Estimation is a valuable procedure to account for undetected evasion, it also has limitations. Most importantly for our purposes in this paper, DCE deals with the possibility of differential undetected evasion across the income distribution only coarsely. The DCE adjustment is applied separately for low- vs. high-visibility income and for taxpayers with reported total positive income above vs. below $100,000. The underlying assumption is that within a visibility $\times$ income bin, undetected evasion is proportional to detected evasion. However, wealthier taxpayers may be more likely to use sophisticated forms of evasion. Thus, undetected evasion may be larger (relative to detected evasion) at the top. Wealthier individuals also have more complex tax returns, making it harder for auditors to uncover all noncompliance.

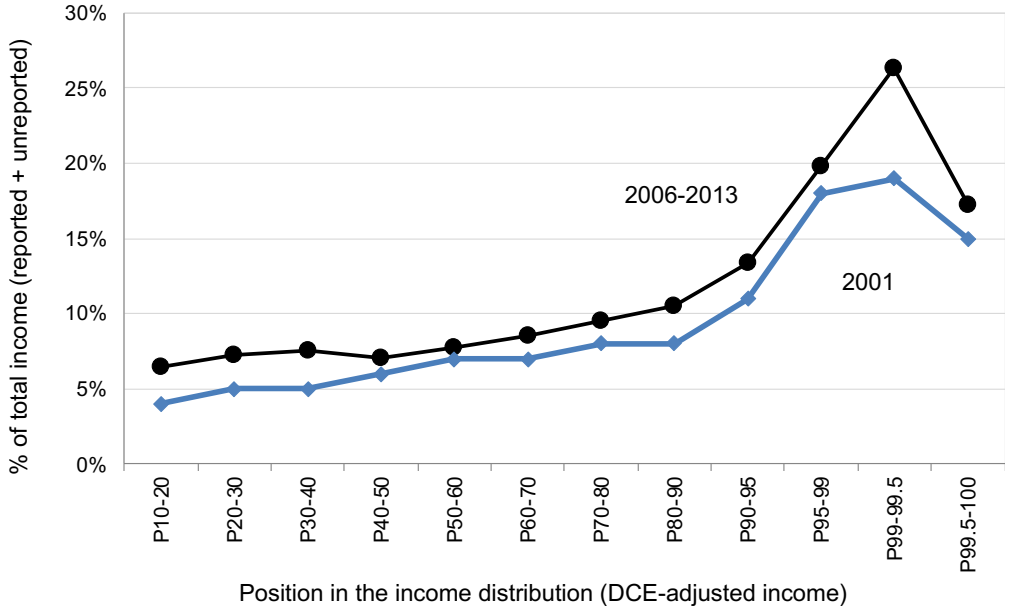
Based on what we know about the information available to auditors and the audit process, there are at least two reasons to believe that the ratio of undetected to detected evasion may rise with income. First, interest, dividends, and capital gains accruing to offshore accounts (which according to available evidence is highly concentrated, e.g., Alstadsaeter et al., 2019) were subject to limited information reporting during our sample period. Second, if a wealthy taxpayer owns a network of private business interests, the auditor faces a considerable challenge in trying to assess the compliance of every single entity in the network. Upon initial review, the auditor checks whether the income allocated to the individual taxpayer by these businesses is accurately reported on the individual tax return, and whether the taxpayer has an active or passive role in the businesses. Internal procedures, the materiality of risk, and the available tools and resources guide the extent to which the broader network is examined. We discuss this further in Section 3.2.

We now provide a simple and direct test of the hypothesis that random audits may miss some top-end evasion, even after the DCE adjustment. We compare the amount of evasion estimated at the top in random audits with the amount found in operational audits. Operational audits include all audits other than random audits: correspondence audits, conventional in-person audits, sophisticated audits of high-income/high-

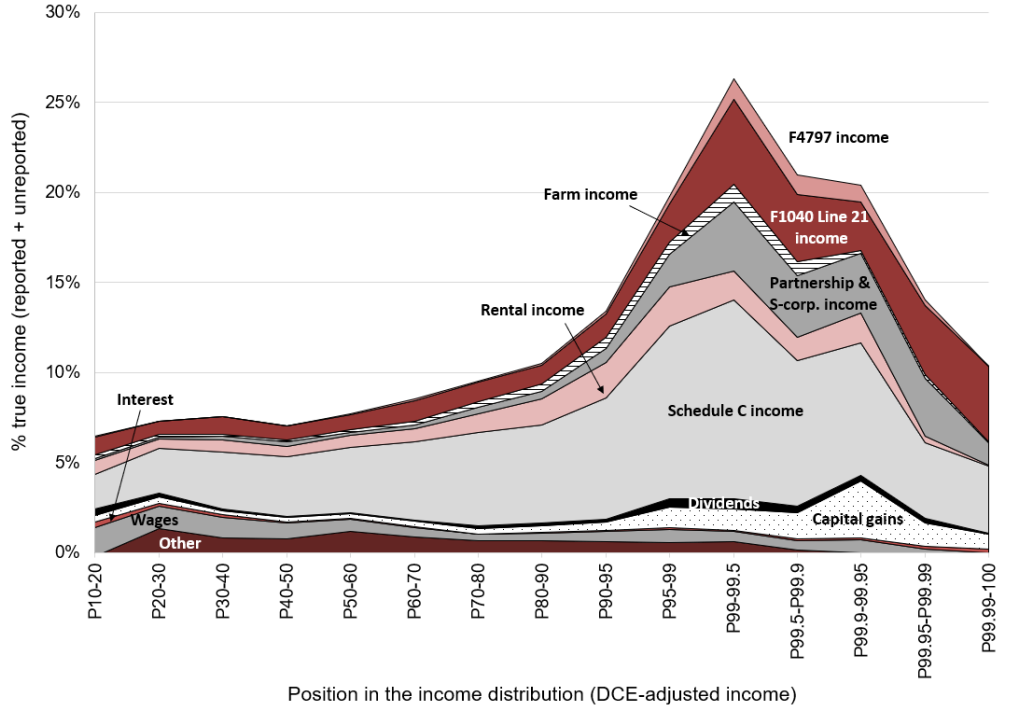
[14] See Appendix Figure A4. One reason for this difference is that we include self-employment tax in our measure of taxes paid, while Johns and Slemrod (2010) do not. Self-employment taxes (15.3% of self-employment income up to the Social Security cap) are large at the bottom of the distribution relative to federal income taxes paid. See also footnote 8.

FIGURE 3: UNREPORTED INCOME IN RANDOM AUDIT DATA AFTER DCE CORRECTION

(a) Unreported Income (% of True Income)



(b) Decomposition by Type of Income (2006–2013)



Note: This figure shows the distribution of under-reported income in the 2006-2013 NRP data with the DCE adjustment. In the top panel we compare our estimates to those in Johns and Slemrod (2010), which are based on the 2001 NRP data and use the same DCE adjustment. Because the top group reported in Johns and Slemrod (2010) is the top 0.5%, we proceed similarly in that panel. In the bottom panel, we show smaller groups at the top (as in Figure 2). Taxpayers are ranked by Adjusted Gross Income (AGI) in Johns and Slemrod (2010) and market income in our series (defined as total income reported on form 1040 minus Social Security benefits, unemployment insurance benefits, alimony, and state refunds), both after DCE adjustment. The difference between these definitions of income is negligible for under-reporting gaps at the top.

wealth individuals, and a variety of other specialized audit programs. Because the IRS prioritizes audits of taxpayers who, based on a variety of factors, it expects to be noncompliant, operational audits are not conducted at random. Because only a fraction of the population is subject to an operational audit in a given year, total evasion in raw operational audit data should always be lower than the population-weighted NRP estimates of total evasion. Appendix Figure A7 depicts the fraction of tax units subject to an operational audit by fiscal year for the top 1%, the top 0.1% and the top 0.01% of reported AGI over time. In 2010 for instance, 10% of tax units in the top 0.01% were audited. Consistent with publicly available data, audit rates increased from the beginning of our sample until about 2013 and then fell due to budget cuts.

Figure 4 contrasts unpaid taxes found in operational audits to the tax gap estimated in the NRP. We focus on fiscal year 2010 because this was the year with the highest total assessed tax in our operational audit data, but data from other years tell a similar story.[15] Because in operational audit data we only observe reported income in the year of the audit (not corrected income), we rank tax units by reported income in Figure 4. We stress an important caveat: when ranking by reported income, it is not possible to make inferences about the true level of tax evasion at the top. Ranked by reported income, top earners by construction tend to have low evasion (since they are selected on high declared income). Figure 4 is thus not informative about the tax gap in the (true) top 0.01%. We only use Figure 4 to compare operational and random audits.

Focusing first on NRP estimates before DCE adjustment, we observe in Figure 4 that although only 10% of taxpayers in the top 0.01% were audited in 2010, assessed noncompliance in operational audits is already much higher than the no-DCE NRP estimate for the entire top 0.01%. In fiscal year 2010, operational auditors assessed $1.5 billion (in 2012 dollars) in additional taxes owed for the top 0.01% by reported income. Assuming that all un-audited taxpayers evaded zero taxes, this implies that the top 0.01% by reported income evaded at least 1.7% of taxes owed.[16] This lower bound for the amount of evasion in the top 0.01% is far larger than the NRP estimate, which is less than 0.1%.[17] In the next two bins—P99.9 to P99.95 and P99.95 to P99.99—the assessed tax gap in operational audits is as large as the population-wide NRP estimate before DCE, despite the fact that less than 10% of taxpayers in these groups were subject to an operational audit. This finding supports the notion that some evasion at the top end is missed by the NRP.

Figure 4 shows that the same observation about the top 0.01% is true of the DCE-adjusted NRP too.[18] When ranking by reported income, there is less evasion in the top 0.01% in the DCE-corrected NRP data than

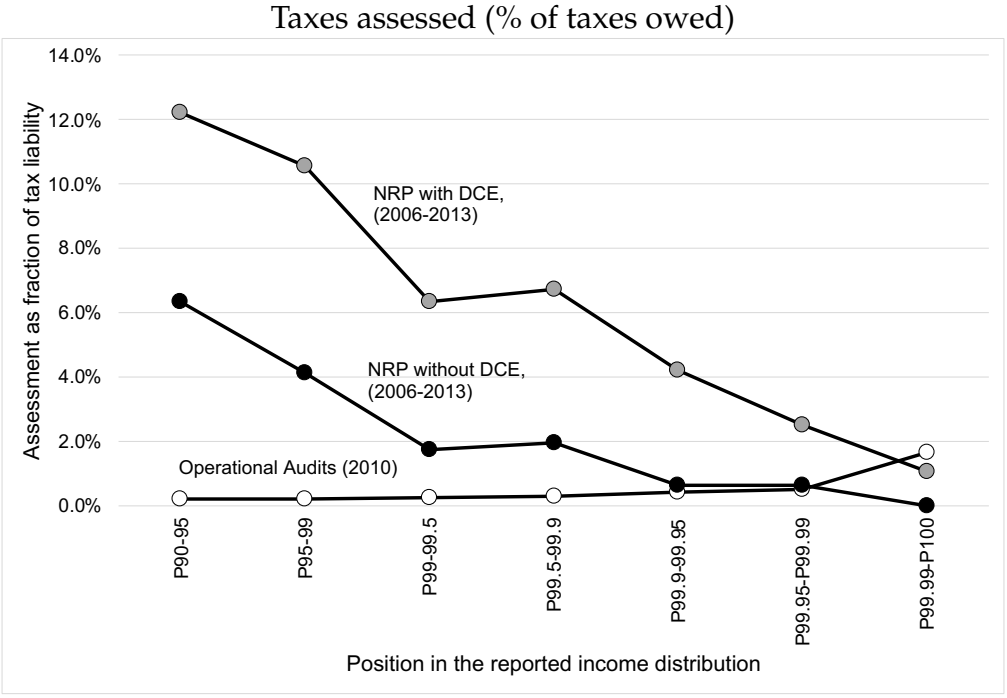
[15] Appendix Figure A8 describes the data plotted in Figure 4 for operational audits of the top 0.01% by fiscal year.

[16] Total reported tax due is about $90 billion (2012 dollars) in this bin.

[17] Ranking by exam-corrected income the corresponding figure is 0.4% of tax owed, which is still well below the operational audit figure.

[18] We note that the difference between the NRP tax gap with DCE and the operational audit tax gap in Figure 4 may not be statistically significant, due to sampling noise in the NRP. However, a statistically significant difference is not necessary for the substantive point we make here, which is that enough evasion is detected in operational audits to imply that the NRP point estimate is inferior to the true tax gap at the top of the income distribution.

FIGURE 4: TAX GAP IN OPERATIONAL AUDITS VS. (POPULATION-WEIGHTED) RANDOM AUDITS



Notes: This figure compares tax evasion detected in operational audit data to population estimated noncompliance in the NRP, focusing on the top 10% of the distribution. Unlike in other figures, we rank individuals by reported income. We plot total tax evaded as a fraction of total tax due in each bin from operational audits in 2010 and compare this to NRP exam corrections with and without DCE. Even with DCE, random audits uncover a very small amount of evasion in the top 0.01% by reported income. Operational audits uncover more evasion than the NRP point estimate in the top 0.01%, even though the operational audit estimate only accounts for evasion by audited taxpayers (about 10% of top 0.01% returns were audited in 2010).

in operational audits. In effect, the DCE correction adds little evasion to detected evasion in the top 0.01% by reported income. Virtually all evasion attributed to the top 0.01% in Figure 3 comes from individuals initially reporting income below the top 0.01% threshold who are re-ranked into the top 0.01% after DCE adjustment (see Figure A3). The DCE methodology itself barely increases noncompliance for the top 0.01% by reported income, which remains below the noncompliance found in operational audits.

In summary, random audits face two main issues. One is that examiners vary in their propensity to detect evasion. The DCE methodology was designed to address this issue. Second, the available tools, procedures, and resources place limits on the extent to which some types of noncompliance at the top can be identified by any examiner conducting a random audit—a limitation inherent in any feasible random audit program. This issue introduces the possibility that even after DCE correction, estimated evasion may be underestimated at the top-end. A direct comparison with operational audit data supports this hypothesis. Until recently, however, there was little direct evidence about the nature and the size of the noncompliance that may go undetected in random audits, making it hard to quantify this issue—a task we now pursue.

3 What Random Audits Miss: Evidence

In this section, we provide a first empirical demonstration that two forms of evasion—the concealment of offshore wealth and tax evasion via pass-through businesses—are (i) highly concentrated at the top of the income distribution, (ii) effectively invisible in random audit data (including after DCE correction), and (iii) quantitatively important for the measurement of the tax gap at the top. For each of these types of evasion, we explain the practical limits faced by random audits, and we show that accounting for these limits implies a large upward adjustment to estimates of noncompliance at the top of the income distribution. To clearly characterize what is detected in random audit data and what is often not detected, in this section we focus on how undetected evasion modifies raw (i.e., non-DCE-corrected) NRP evasion. We proceed separately for offshore and pass-through business evasion. In section 4, we will combine offshore evasion, pass-through evasion, and the DCE adjustment to construct our preferred estimate of the level and distribution of federal income tax noncompliance in the United States.

3.1 Offshore Evasion

We start by showing that evasion conducted through offshore financial accounts is highly concentrated at the top of the income distribution and almost never detected by NRP auditors in the period preceding increased transparency in offshore reporting and enforcement. Our benchmark year in this sub-section is 2007, i.e., the year preceding the start of a series of initiatives to fight offshore tax evasion. We take the pooled 2006–2013 NRP data as representative of the year 2007, and, leveraging the retrospective information created by the post-2007 crackdown, ask how accounting for offshore evasion modifies the level and distribution of detected evasion in 2007. We also discuss how these results can inform knowledge about top-end evasion post-crackdown.

3.1.1 Background and Data on Offshore Evasion

In 2008–2009, the IRS and the U.S. Justice Department began an ambitious crackdown on offshore tax evasion, described in more detail in Johannesen et al. (2020). Key steps in this process included the establishment, starting in 2008, of Offshore Voluntary Disclosure (OVD) programs whereby taxpayers could disclose prior noncompliance and pay penalties but avoid potential criminal prosecution, the passage in 2010 of the Foreign Accounts Tax Compliance Act, and the implementation of FATCA third-party reporting for offshore accounts in 2014. Johannesen et al. (2020) find that enforcement caused a large increase in reporting of offshore wealth and the associated financial income by U.S. taxpayers. We build on these findings to construct two datasets of individuals that are very likely to have been evading taxes on income from their offshore

assets prior to the crackdown. The data in this paper are the same data used in Johannesen et al. (2020), slightly updated to include additional years of the OVD program.

The first dataset of likely evaders are participants in the Offshore Voluntary Disclosure Program. We gathered data on all participants in OVD Programs from 2009 to 2015 and matched 50,020 OVD participants[19] to their individual tax returns. We refer to this sample as the *OVDP participant* sample.

The second dataset of likely evaders we use consists of individuals reporting that they own offshore assets by filing a Foreign Bank Account Report (FBAR) for the first time between 2009 and 2011. U.S. persons that are the beneficial owners of more than $10,000 in offshore financial wealth have been required to disclose this wealth to the government since the 1970s by filing an FBAR. We use only those first-time FBAR filers with U.S. addresses, disclosing an account in a tax haven.[20] Johannesen et al. (2020) found compelling evidence that the large majority of these taxpayers had been evading U.S. tax on these assets prior to disclosing them in response to enforcement. We match 31,752 such taxpayers to their individual income tax returns. We note that this sample may contain individuals who had an offshore account for legitimate reasons and were unaware of their FBAR filing obligation prior to increased enforcement, but we find it less plausible that individuals at the very top of the income distribution with accounts in havens were unaware of these obligations. We refer to this sample as the *first-time FBAR filer* sample.

For both sets of taxpayers, we then use data from their income tax returns to rank them in the income distribution. Specifically, we use income data for the tax year after these individuals' disclosure of their offshore wealth, as the results in Johannesen et al. (2020) suggest that this is the year in which individuals start to comply with their tax reporting obligations on such wealth and associated income. We rank individuals by adjusted gross income (AGI) for simplicity, but we show that one obtains similar results with alternative rankings.

For the first-time FBAR filer sample, we also use data on the amount of offshore wealth disclosed on their FBARs. These particular FBAR filers, those with U.S. addresses newly disclosing tax haven accounts, disclosed $124 billion in wealth between 2009 and 2011. For comparison, total reported FBAR wealth was about $290 billion for a given year in the same period, suggesting that a sizable share of the overall wealth reported on FBARs in this period came from newly disclosures of wealth in tax havens from filers with U.S. addresses. However, estimates of total wealth concealed in tax havens are higher. As we discuss below, our preferred estimate of this amount (form Alstadaeter et al., 2018) is just over $1 trillion, suggesting that

[19]This 50,020 figure does not include approximately 6,000 program participants who we could not match to their individual tax returns. About two-thirds of these were businesses participating in the OVD program. The rest did not file a tax return in the year we wished to analyze, either because their participation in OVDP was too recent or for some other reason.

[20]We use the same list of tax havens as Johannesen et al. (2020), which is the OECD (2000) list of uncooperative tax havens plus Switzerland, Hong Kong, Singapore, and Luxembourg. We note that this is not an official definition used by the IRS, which has no official list of countries it considers a tax haven.

roughly 11 percent of all offshore wealth in tax havens is disclosed by individuals in our first-time FBAR filer sample. All of this is consistent with the results of Johannesen et al. (2020).

3.1.2 Empirical Analysis of Offshore Evasion

We first show that NRP audits very seldom detected offshore evasion in the time period we study. As both of our offshore disclosure samples and the NRP stratified random audit sample contain disproportionately many observations of high-income individuals, it turns out that there is enough overlap between them for a simple, direct statistical analysis. Specifically, 378 first-time FBAR filers and 135 OVDP participants were selected for NRP audits between 2006 and their disclosure of an offshore account. In Figure 5a, we show that the auditor discovered that the individual had offshore wealth and should have been filing an FBAR in only 7 percent of these cases. It is likely that the vast majority of these taxpayers had offshore wealth in the period immediately preceding their disclosure—offshore wealth that they should have disclosed in the year they were audited.[21] This finding suggests that NRP random audits seldom detected concealed offshore wealth.

Figure 5b depicts the fraction of the overall population in a particular range of income that are present in one of these samples of likely evaders, accounting for the overlap between samples in the total (as OVDP participants were required to file any delinquent FBARs as part of participating in the program). We find a very steep profile, with those at the very top of the income distribution being significantly more likely to appear in the samples of likely evaders, even within the top 0.1%. The profile is slightly less steep for OVDP participants than for other first-time FBAR filers. Altogether, almost 7% of taxpayers in the top 0.01% of the income distribution (999 taxpayers) were part of one of these samples of likely evaders.[22]

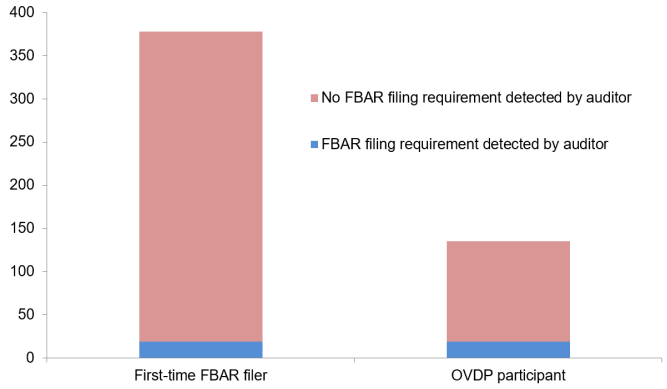
In Figure 5c, we turn to the distribution of wealth by rank in the income distribution. For this analysis we used only the first-time FBAR filer sample, where we have an arguably good measure of offshore wealth from the FBAR itself. We calculate the share of all wealth reported on FBARs in this sample that is attributable to taxpayers at different parts of the income distribution. For contrast, we use estimates from

[21] Churn in the population of owners of offshore accounts could imply that some of the individuals disclosing offshore wealth in our data did not own an offshore account when they were audited under the NRP, but this possibility seems unlikely to affect our key takeaways. In a given year, about 30% of all FBAR filers do not file in the subsequent year; turnover is smaller for large accounts in tax havens. Even if we supposed that only 70% of our overlap sample actually owned an offshore account when they were audited under the NRP, we would conclude that the auditor only detected offshore wealth in $0.07/0.7 = 10\%$ of cases. In other words, the overall detection rate is so low that no realistic amount of churn will imply a high detection rate. Relatedly, in most of the 7% of cases where the auditor did indicate that the individual had an FBAR filing requirement, they found that the individual had complied with their filing requirement, suggesting that no noncompliance was found even in many of these cases.

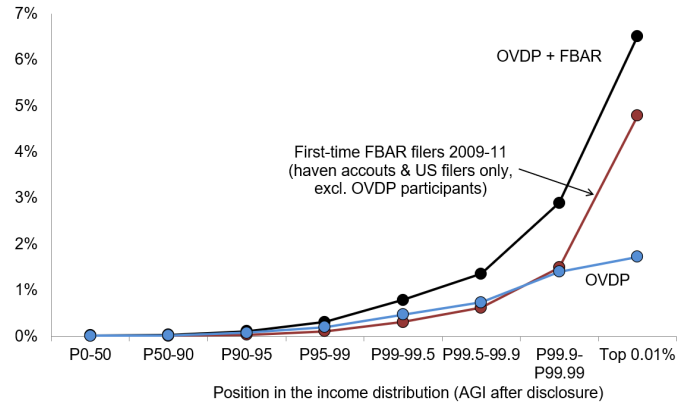
[22] Appendix Figure A5 shows that our choice of income definition matters little for these estimates. We plot the fraction of the population in the first-time FBAR filer sample by rank in the income distribution for various definitions of income. We start with Adjusted Gross Income from the tax return, as in Figure 5b. As we observe some individuals with large business losses holding substantial offshore wealth, we also rank individuals by total positive income, which replaces the income components of AGI that can be negative—net capital losses and business losses—by zero when they are negative. Finally, we rank people by realized financial capital income, defined as the sum of interest, dividends, and realized capital gains and losses. The overall profile is very similar for the three different income concepts, though it is steepest for financial capital income, followed by positive income.

FIGURE 5: FINDINGS FROM DATA ON OFFSHORE EVASION

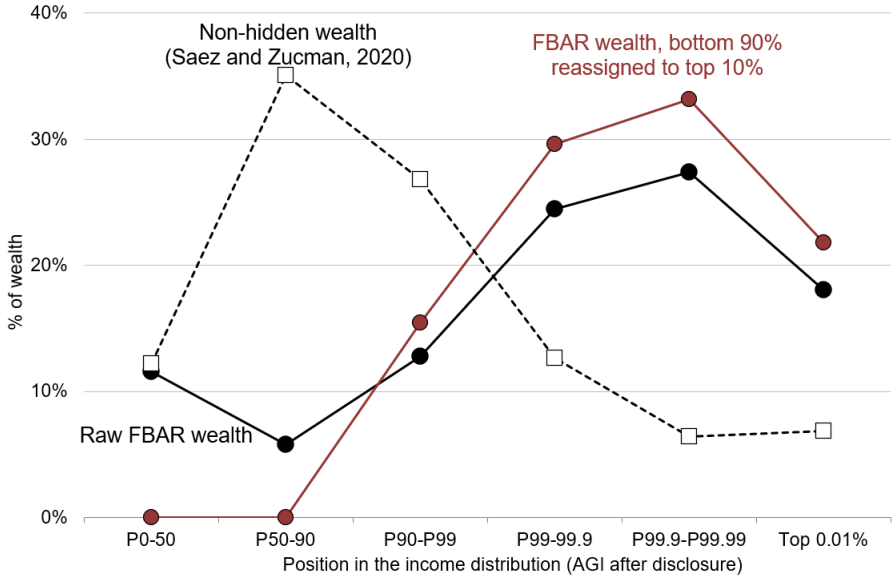
(a) Do NRP audits detect offshore evasion?



(b) Share Disclosing by Income Rank



(c) Distribution of Offshore Wealth



Note: This figure summarizes the main findings from data on the OVDP participant and first-time FBAR filer samples of likely evaders with respect to offshore wealth. Panel (a) shows that for both samples, individuals within the sample that happened to be audited in the NRP were virtually never discovered (see also footnote 21). These individuals nevertheless disclosed an offshore account in a later year. Panel (b) plots the fraction of the full population in each sample by bin of adjusted gross income (in the tax year after disclosure of the offshore account), accounting for overlap between the samples. We observe the steep profile of the probability of disclosing a previously hidden account by income rank. The main difference between the profiles from the two samples appears to be the presence of many more individuals in the top 0.01 % of the income distribution in the first-time FBAR filer sample. In total nearly 7% of people in the top 0.01 percent of the income distribution appears in one of the two samples. Panel (c) plots wealth shares for non-hidden wealth from Saez and Zucman (2016) updated in Saez and Zucman (2020) by bin of market income (defined as total income reported on form 1040 minus Social Security benefits, unemployment insurance benefits, alimony, state refunds, and other income), versus wealth reported on FBARs by the first-time FBAR filers with U.S. addresses and accounts in tax havens (ranking by positive market income after disclosure). We observe that FBAR wealth is extremely concentrated at the top of the income distribution.

the capitalization method of Saez and Zucman (2016) updated in Saez and Zucman (2020) to depict wealth shares for non-hidden wealth.[23] We observe that FBAR wealth held in havens is much more concentrated at the top of the income distribution than non-hidden wealth, with over 20% of the wealth belonging to the top 0.01% by income, as opposed to 7% of non-hidden wealth. For reference, the total population of FBAR filers discloses $124 billion in offshore wealth, with about $26 billion in the top 0.01% and $36 billion between the 99.9th and 99.99th percentiles. This result confirms that the findings in the previous figures were not simply driven by the overall concentration of wealth at the top of the income distribution, but rather that concealed offshore wealth is *especially* concentrated at the top.

Further analysis suggests that even the modest amount of FBAR wealth attributed to the bottom 90% of the income distribution may actually belong at the top of the distribution. Most of the observed FBAR wealth in the bottom 90 percent of the income distribution is driven by a small number of extremely large accounts, which skews the distribution of FBAR wealth in these income groups. For instance, the median level of FBAR wealth for those in the bottom 50% of the income distribution is around $200,000 and the mean is $2.5 million. Thus a small number of high-wealth individuals in the bottom 90% accounts for almost all of the offshore wealth in this group. We suspect that the vast majority of the FBAR wealth attributed to the bottom 90% of the income distribution (11% of all FBAR wealth using total positive income and 17% using AGI) should in fact be assigned to top income groups, and would be assigned to the top if we ranked by wealth instead of income. To illustrate how much this matters, in Figure 5c we depict the impact on the FBAR wealth shares of reassigning wealth from the bottom 90 percent to the top 10 percent of the income distribution, in proportion with the FBAR wealth already attributed to the top 10%.

A caveat to our analysis of offshore wealth deserves to be noted. Both of our samples contain data on offshore wealth of *voluntary* disclosers of offshore wealth, those who selected to participate in the OVDP or to engage in a likely quiet disclosure (see Johannesen et al., 2020).[24] As such they cannot be regarded as a representative sample of all owners of offshore wealth. To more fully understand the distribution of offshore wealth, we would ideally combine these data with samples closer to a random draw from the population of offshore evaders. We do not have direct evidence on the direction of the selection, but data from other countries suggests offshore wealth may be even more concentrated at the very top than we estimate. Alstadsaeter et al. (2019) use leaked data from HSBC Switzerland and estimate that 52% of offshore wealth in that bank was owned by taxpayers in the top 0.01 percent of the wealth distribution in Scandinavia.

[23] These series incorporate valuable suggestions made by Smith et al. (2019) and Auten and Splinter (2019). As our focus is documenting the large difference between ownership shares of offshore wealth and those of domestic wealth rather than comparing the evolution of wealth shares over time, using any alternative estimates of wealth shares would lead to similar findings.

[24] A "quiet disclosure" is when a taxpayer begins to report a previously undisclosed foreign account and the income in that account without participating in an OVD program.

TABLE 1: OFFSHORE EVASION SCENARIOS

Parameter	Lower-bound scenario	Preferred scenario	Upper-bound scenario
Amount of U.S. offshore wealth (in billion $)	750	1,058	1,500
Fraction of offshore wealth concealed	85%	95%	100%
Rate of return on offshore wealth	4.65 %	6%	11%
Distribution of offshore wealth	FBAR	Average of FBAR and Nordic	Nordic
Average Marginal Tax Rate	20%	25%	30%

Note: This table summarizes the five sets of assumptions about the amount and distribution of offshore income made in our three different scenarios discussed in Sections 3.1.3 and 3.2.4.

3.1.3 Implications of Undetected Offshore Income for the Distribution of Noncompliance

We now consider the magnitude of the adjustment to the income under-reporting gap implied by accounting for undetected offshore evasion. Following Alstadsaeter et al. (2019), we obtain an estimate of the amount of unreported offshore income by proceeding in four steps. Each step entails an assumption, which we list in the middle column of Table 1. We modify these assumptions in subsequent sensitivity analysis (see the first and third columns of Table 1) to quantify the margin of error involved in our adjustment.

In our preferred scenario we make the following assumptions. First, we start with an estimate of aggregate offshore wealth in tax havens owned by U.S. households in 2007: $1,058 billion, the equivalent of 1.7% of total U.S. household wealth. This number is taken from Alstadaeter et al. (2018), Appendix Table A.3, with no modification; it was obtained by Alstadaeter et al. (2018) by allocating the 2007 global amount of offshore wealth estimated in Zucman (2013) to each country, using retrospective statistics on the ownership of offshore bank deposits released by the Bank for International Settlements in 2016. Second, we assume that 95% of that wealth was hidden. Some accounts were certainly properly declared in 2007, but a 95% rate is consistent with the United States Senate (2008, 2014) reports, which found that 90%–95% of the wealth held by American clients of a number of Swiss banks were undeclared before FATCA.

Third, we assume a nominal taxable rate of return on this offshore wealth of 6.0%. This rate of return is inferred from what is known about the portfolio composition of global offshore wealth around 2007 and the rate of return on assets at that time. More precisely, Zucman (2013) estimates that in 2007, around 75% of global offshore wealth was invested in securities (mostly equities and mutual fund shares) and 25% in bank deposits. Our 6% taxable return is obtained by assigning the average interest rate paid by Swiss banks to deposits, and assigning half of the S&P 500 return to securities (with the other half consisting of unrealized capital gains).[25]

[25]The average interest rate paid by Swiss banks on their term deposits was 4.3% in 2006 (the U.S. Federal fund rate was in range of 4.3% to 5.25%). The total nominal return (dividends reinvested) was 13.4% for the the S&P 500 (and 20.65% for the MSCI world). With 25% of assets earning a 4.3% return in bank deposits and 75% earning half of a 13.4% return in securities (with the other half in unrealized capital gains), we arrive at our taxable 6% return.

Fourth, we distribute the macro amount of offshore wealth and income as follows: We take a weighted combination of the distribution of offshore wealth observed among self-selected U.S. filers who disclose a haven account by filing an FBAR for the first time in 2009-2011 (depicted in Figure 5c), and the distribution of hidden wealth estimated by Alstadsaeter et al. (2019) in Scandinavia. We put equal weight on the U.S. disclosed offshore wealth distribution and the Alstadsaeter et al. (2019) distributions. This implies that 60% of hidden wealth belongs to the top 0.1% highest earners and 35% to the top 0.01% (vs. more than 50% in Scandinavia). Ideally, of course it would be preferable to base our allocation of offshore wealth only on U.S. data. We hope to refine our allocation in the future using additional U.S. data where self-selection might be more limited.

Under these assumptions, unreported offshore income adds up to 0.7% of aggregate taxable income in 2007. As we saw in section 2.1, before any correction for undetected evasion (in particular before DCE), the NRP finds that 3.9% of income is under-reported. Adding unreported offshore income increases this number to 4.5%. Figure 6a shows how adding offshore income modifies the distribution of noncompliance. Unsurprisingly, adding offshore income has no visible effect in the bottom 90% of the distribution and only a small effect between the 90^{th} and 99^{th} percentile. However, although offshore evasion is small on aggregate, accounting for it makes a significant difference at the top. It increases the ratio of under-reported income to true income by 4.0 percentage points in the top 0.01%, and by 2.9 percentage points for the top 0.1% excluding the top 0.01%. As a result, the sharp drop-off in the income under-reporting and tax gap by income within the top 1% is undone by accounting for offshore evasion alone.

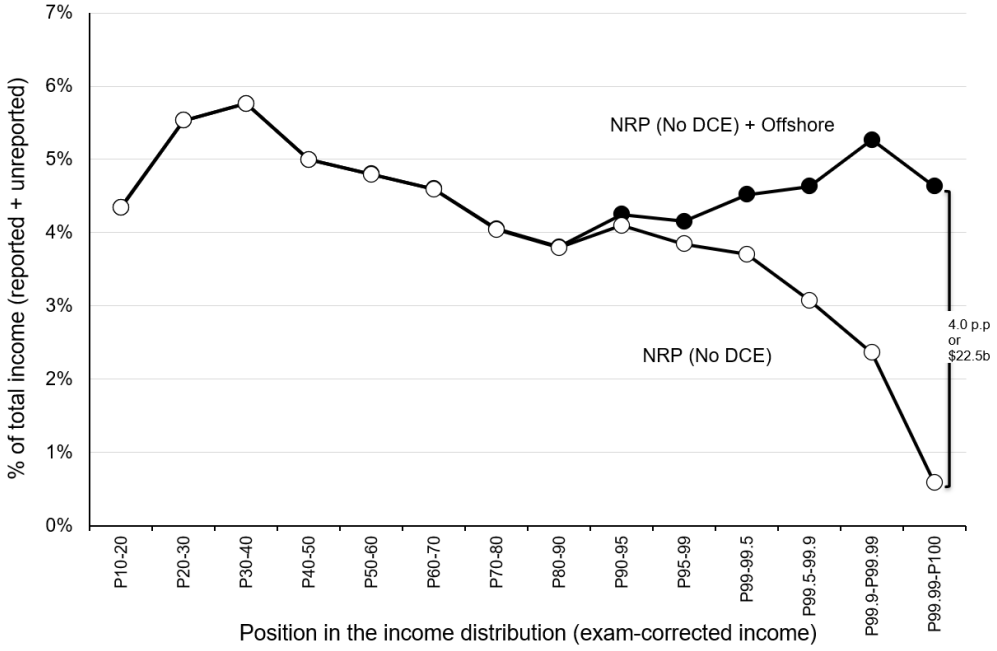
To estimate federal income tax evaded as a share of tax due, we must make a fifth assumption about the average marginal tax rate on income from offshore wealth. The average marginal tax rate should be between the marginal tax rate on ordinary income and the preferred tax rate on long-term capital gains and qualified dividends, which are 35% and 15% in our reference year, respectively. Reflecting our earlier discussion about the portfolio composition of offshore wealth, we use an average marginal tax rate of 25% in our preferred scenario.[26] Appendix Figure A10 presents estimates of the amount of taxes evaded as a fraction of taxes owed, adding evasion on offshore wealth to the raw NRP estimates before DCE correction. In total, we estimate that $15 billion in taxes was evaded from offshore accounts, with $10.5 billion of this total attributed to the top 0.1%, and $6.4 billion attributed to the top 0.01%.[27] Accounting for offshore

[26] This rate is consistent with a scenario in which 25% of taxable offshore income is interest income, 50% is long-term capital gains and qualified dividends, and 25% is short-term capital gains and non-qualified dividends; see footnote 25. We provide sensitivity checks for a 20% and 30% tax rate in Section 3.2.4.

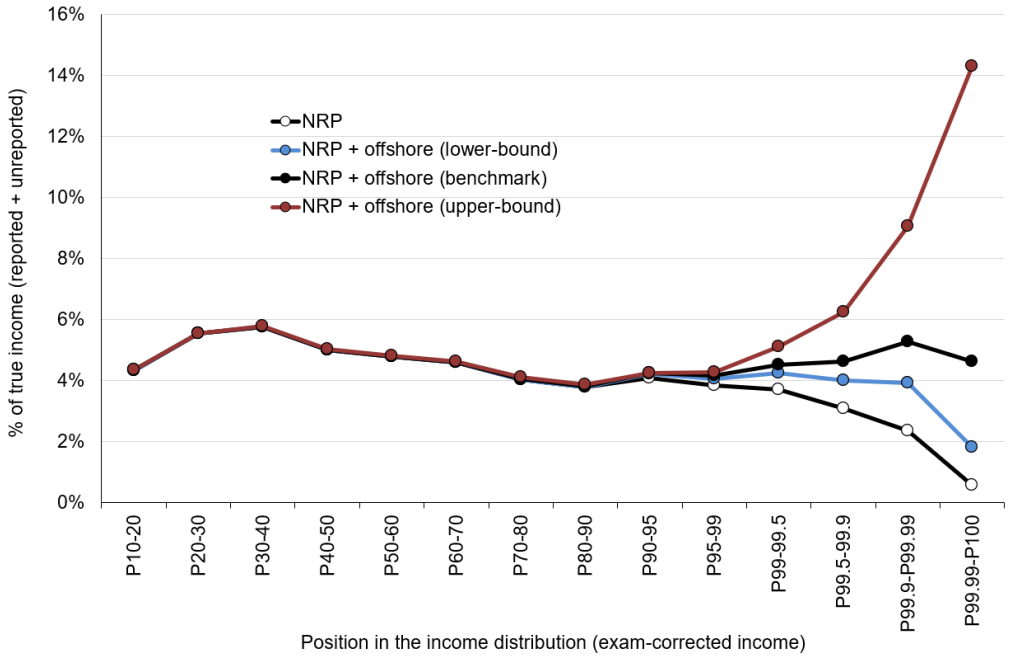
[27] Our benchmark estimate of $15 billion in revenue loss is lower than the $23 billion estimate in Zucman (2014), because Zucman (2014) includes both federal and state taxes (and thus applies a 30% combined federal-plus-state marginal tax rate, as opposed to 25% in our benchmark scenario that captures federal taxes only) and assumes a 7% return (vs. 6% in our benchmark scenario). Zucman (2014) also estimates evaded estates tax (assuming 3% of offshore wealth belongs to decedents and a 40% estates tax rate), leading to total (income plus estate, federal plus state) tax evaded on offshore wealth of $36 billion (Zucman, 2014, Table 1 p. 140).

FIGURE 6: ACCOUNTING FOR UNDETECTED OFFSHORE FINANCIAL INCOME

(a) Unreported Income (% True Income)



(b) Sensitivity Analysis



Note: This figure plots the estimated income under-reporting rates with and without adding offshore tax evasion. The top panel shows our preferred scenario and the bottom panel reports our sensitivity analysis. Taxpayers are ranked by exam-corrected market income in the NRP data, and offshore adjustments are made on the basis of positive market income; this is the best available estimate of "true income" before DCE adjustments. We find that income under-reporting rates increase significantly at the top of the income distribution when accounting for offshore evasion, reversing the sharp drop-off in estimated evasion at the top seen in uncorrected random audit data. The point estimate for the top 0.01 percent increases by 4 percentage points in our benchmark scenario.

evasion increases the tax gap at the top of the distribution significantly.

It is worth noting that accounting for offshore evasion also implies an upward revision of DCE-adjusted estimates of under-reported income and unpaid taxes at the top, because the DCE adjustment is unlikely to capture tax evasion via offshore financial assets. As shown by Appendix Table A1, a very small fraction of dividends and interest are found to be under-reported in the NRP, 2.8% and 1.4% respectively—or 0.13% of total taxable income. In the DCE methodology, interest and dividends are grouped with other "high-visibility" income sources such as wages (IRS, 2019, Table 4 p.18) because they are subject to third-party reporting in the U.S., and therefore get a low multiplier (of about two). As a result, the DCE methodology increases detected evasion on interest and dividends by around 0.1% of total taxable income. This is much less than our benchmark estimate of unreported offshore income, 0.7% of total income.

3.1.4 Offshore Tax Evasion: Sensitivity Analysis

We present results of our sensitivity analysis in Figure 6b. For simplicity, we focus on two scenarios: one in which each of our assumptions is chosen (given the available evidence) to minimize the amount of offshore evasion at the very top, and one in which each assumption is chosen to maximize it (again given the available evidence). These scenarios provide plausible lower and upper bounds for the size of offshore tax evasion at the top of the income distribution. The first and third columns of Table 1 describe these alternate assumptions.

The lower bound of the amount of offshore wealth ($750 billion) comes from the Boston Consulting Group's Wealth Report of 2007, which estimated that wealthy North American residents held about $37.7 trillion of wealth, 2% of which was held offshore.[28] The upper-bound is based on Guttentag and Avi-Yonah (2005), who built on the BCG Wealth Report of 2003, according to which the total holdings of high-net-worth individuals in the world were $38 trillion, including $16.2 trillion for North America residents. "Less than 10%" of this wealth was held offshore according to the BCG; using this percentage as an upper bound as in Guttentag and Avi-Yonah (2005) gives an approximate $ 1.5 trillion of U.S. offshore wealth. The lower-bound of the fraction of offshore wealth which is hidden is based on United States Senate (2008, 2014) reports investigating the practices of several Swiss banks in the U.S. In these reports, the investigation committee find that about 90% of the wealth held by U.S. taxpayers at UBS Switzerland was undeclared, and that between 85% and 95% of the accounts held by U.S. taxpayers at Credit Suisse were undeclared. For the rate of return on wealth held offshore, the conservative figure corresponds to the average daily 10-year Treasury rate for the year 2007, while the upper-bound number is the return on average equity for all U.S.

[28] Alstadsaeter et al. (2019, footnote 28 p. 2090) list all the available estimates of the global amount of offshore around 2007; the BCG estimate is the second lowest one, immediately after (and close to) an OECD estimate which is not broken down by country.

banks, averaged over the year 2007. Total income under-reporting via offshore accounts is $60.3 billion in the preferred scenario (0.7% of true total taxable income), $28.7 billion in the lower bound scenario (0.3% of total income), and $165 billion in the upper bound scenario (1.9% of total income). Finally, our preferred estimate of the distribution of offshore wealth and income was a weighted combination of our FBAR distribution and the distribution from leaks in the Nordic countries (Alstadsaeter et al., 2019). For the sensitivity analysis we put 100% of the weight on one or the other of these.

Two conclusions emerge from Figure 6b. First, there is some uncertainty in estimates of unreported offshore income, which is reflected in the margin between the lower-bound and the upper-bound aggregates above. In the upper bound scenario, under-reported income as a share of true income is 9.6 percentage points higher than in our preferred scenario for the top 0.01%, while in the lower bound scenario is it 2.8 percentage points lower. A similar band is found for tax evaded; see Appendix Figure Figure A11 and Appendix Table A3. Second, and interestingly, even in the lower-bound scenario in which concealed offshore income is very small on aggregate (0.3% of taxable income), accounting for offshore evasion still has a large impact on estimated evasion at the top. It erases the downward-sloping profile of unreported income (as a fraction of true income) seen in non-DCE corrected NRP data from the 99^{th} percentile to the 99.99^{th} percentile, although a drop-off remains in the top 0.01% in the lower bound scenario. In the lower bound scenario, accounting for offshore evasion also doubles the amount of tax evasion detected in the NRP for the top 0.01%. The striking and non-obvious result of our computations is that even under very conservative assumptions about offshore evasion, taking this form of noncompliance into account implies large adjustments to detected evasion at the top.

In the Appendix, we unpack each step of the sensitivity analysis to see which assumptions matter most. Appendix Figure A6a builds up the upper bound scenario by modifying assumptions from the preferred scenario one by one; Figure A6b does the same thing for the lower bound scenario. We observe that the taxable rate of return on offshore wealth, especially at the very top of the distribution, is the most important source of uncertainty. Our own assessment is that the low rate of return used in the lower-bound scenario (4.5%, the 10-year Treasury yield) is likely too low for individuals in the top 0.01% of the income distribution in 2007, given that a large fraction of offshore wealth was invested in equities. However, direct evidence on this question is limited. The next-most important assumption after the rate of return is the distribution of offshore assets. Changing this distribution primarily affects the amount of evasion allocated to the 0.01% versus the rest of the top 1%.

Finally, it is worth asking how our results, which are for the year 2007 (before the increase in enforcement effort on offshore wealth) can inform knowledge about top-end evasion post-crackdown. The available evidence suggests that post-2007 enforcement may have substantially reduced offshore evasion (Johannesen

et al., 2020; De Simone et al., 2020). In particular, the implementation of the Foreign Account Tax Compliance Act in 2014 has significantly increased the information available to the IRS. We will take these facts into account in Section 4 when we present our preferred estimates of top-end evasion in the United States. In any case, we view the results in this section as highly informative for the analysis of top-end evasion post-crackdown. These results are the first empirical demonstration in the U.S. context that some forms of evasion are highly concentrated at the top of the income distribution, effectively invisible in random audit data, unlikely to be picked up by the DCE adjustment, and quantitatively important for the overall tax gap at the top. In what follows we provide evidence that another (possibly rising) form of evasion shares these properties, namely tax evasion occurring via pass-through businesses.

3.2 Evasion on Pass-Through Business Income

Pass-through businesses (S-corporations and partnerships) are business that are not subject to the corporate income tax. Instead, all of the income of these businesses "flows through" to their owners' tax returns, where it is subject to tax. Because ownership of pass-through entities is highly concentrated among the highest-income taxpayers, and because the use of pass-through business structures has been on the rise since 1986, obtaining accurate estimates of noncompliance in such structures is increasingly important. In this section we attempt to make progress on this question by leveraging new data, focusing on the benchmark year 2007 to facilitate aggregation (in Section 4 below) with the offshore evasion results.

3.2.1 Background on Pass-Through Businesses

Administratively, pass-through businesses file tax returns reporting entity-level income and its components. Then, they report the income allocated to their owners on the Schedule K-1. When the owner of a pass-through business is active in the management or control of her business, the Schedule K-1 issued by the business is not strictly from a third-party; in businesses where there is a single owner, the potential scope for noncompliance is similar to sole proprietorship income.[29] When the owner is passive, the business may still evade tax and as a result allocate too little income to its owners. Confirming the flow of income, deductions, credits and other tax features of pass-through businesses takes expertise and resources. Partnerships create a specific additional challenge to the audit process, because partnerships can be owned by other entities, sometimes leading to complex ownership structures involving numerous partnerships, corporations, trusts, or other intermediaries (Cooper et al., 2016).

[29] In 2017, according to SOI tabulations of S-corporation tax returns, 66% of all S-corporations (earning 42% of all S-corporation business income) had a single shareholder. When there is more than one owner, misreporting of pass-through income may be riskier than misreporting of sole proprietorship income (Kleven et al., 2016).

In the context of the NRP, resource constraints and the tools available to auditors limit the comprehensiveness of the examination of pass-through income. First, when individuals report partnership or S-corporation income on their individual tax returns, auditors rarely examine the tax returns of the corresponding pass-through businesses. For individuals subject to an NRP audit and reporting pass-through business income, we observe that under-reporting is detected via an audit of an affiliated pass-through business in only 3.8% of cases.[30] This is because significant resources and expertise are necessary to audit up through networks of pass-through returns and ensure that the correct income (and other tax features) flow through to the relevant individual income tax returns. Such situations are often handled through specialized audit programs that are not part of the NRP. Second, due to resource constraints there is no recurring random audit program of pass-through businesses. The most recent random audit program for partnerships was conducted in 1982. A small-scale pilot random audit program of S-corporations was conducted in 2004.

This situation has a number of implications. First, random audit data likely underestimate the overall amount of tax evasion occurring via pass-through entities. In Appendix Table A1, we see that NRP audits (before DCE correction) find that only 4.2% for partnership and S-corporation income is under-reported. This is much below the estimated 36.7% under-reporting rate for sole proprietorship income, which is subject to direct and comprehensive examination in NRP audits. To better understand this result, Figure 7a contrasts the probabilities that noncompliance was detected in NRP random audit for sole proprietorship income (Schedule C) and pass-through business income (Schedule E). About 60% of audited taxpayers with Schedule C income are found to be under-reporting their Schedule C income. By contrast, only 14.5% of taxpayers with pass-through business income are found to be under-reporting their pass-through business income. These cases correspond to cases in which there is a mismatch between business income reported on the Schedule K-1 and income reported on the Form 1040 (or to cases in which active business income is declared as passive, or vice-versa), not to cases where evasion is detected at the business level. Because businesses are rarely audited in the context of the NRP, only 3.8% of taxpayers with pass-through business income are found to be under-reporting income via an audit of the pass-through business itself.[31]

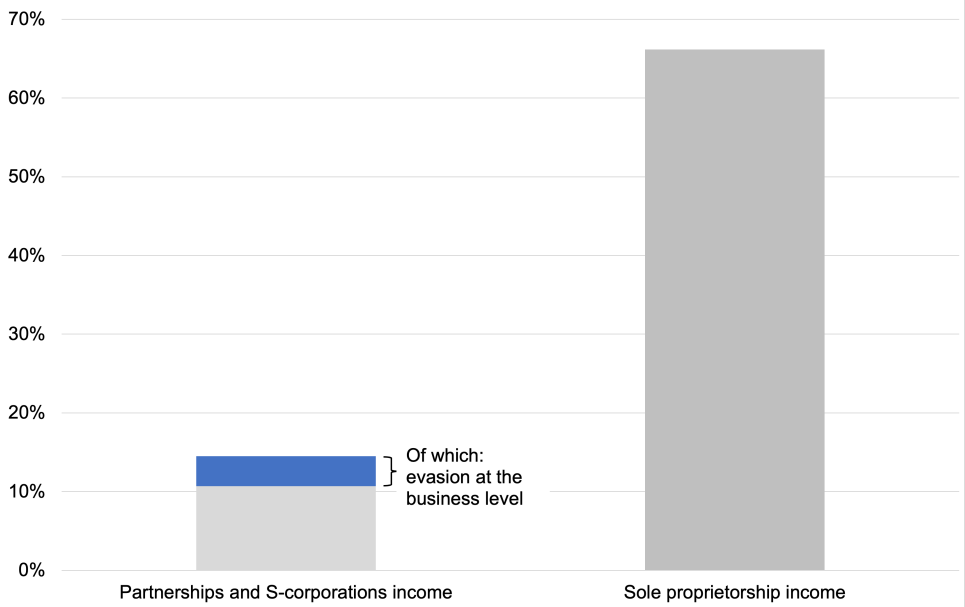
Second, because the ownership of pass-through businesses is concentrated, the bias occurs primarily at the top of the distribution. Figure 7b shows the fraction of income that the various groups of the income distribution earn via pass-through businesses in 2007, our benchmark year. For the bottom 90% of the income distribution, less than 5% of income derives from pass-throughs. However, among the highest

[30] Where audits of affiliated pass-through businesses do occur, the audited pass-through entities are typically small businesses where the individual taxpayer being audited under the NRP likely has access to the partnership books and records.

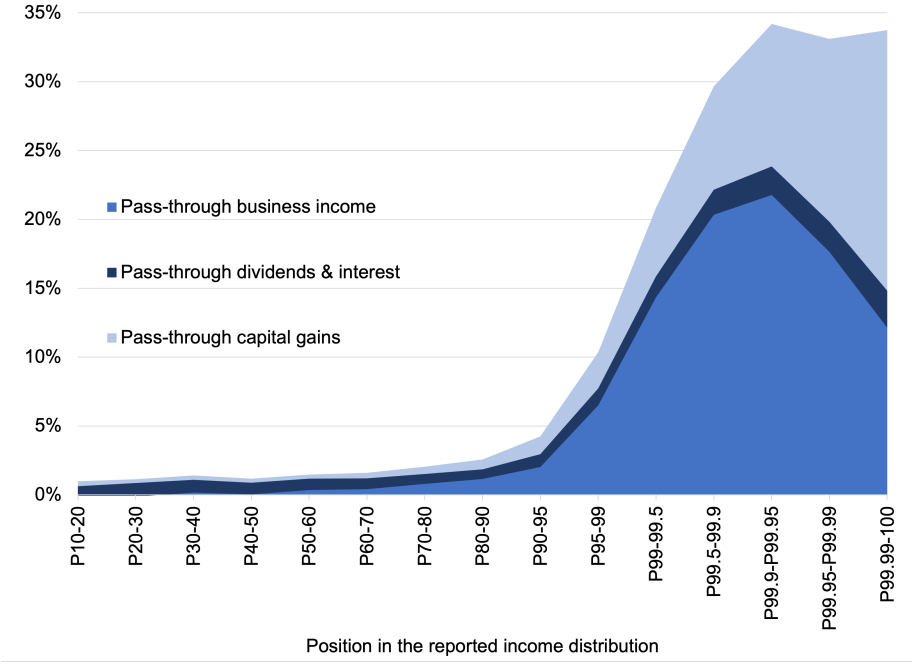
[31] Substantially more NRP related business audits occur for S-corporations than for partnerships. About 75% of the audited related businesses are S-corporation despite the fact that more NRP participants are partnership owners than are S-corporation owners.

FIGURE 7: PASS-THROUGH BUSINESS INCOME: DETECTED EVASION AND CONCENTRATION

(a) Probability to Detect Unreported Business Income in the NRP



(b) Fraction of Reported Income Earned via Pass-Throughs



Note: The top panel shows the probabilities that noncompliance is detected for sole proprietorship income (schedule C) and pass-through business income (schedule E) in 2006–2013 NRP random audits, before any DCE adjustment. Only 14.5% of taxpayers with pass-through business income are found to be under-reporting their pass-through business income. Among these cases, only a quarter correspond to cases in which evasion is detected at the business level. The bottom panel shows the fraction of reported market income earned via pass-through businesses by reported market income in 2007. Income earned via pass-through businesses includes business income (S-corporation and partnership profits) and investment income (capital gains, dividends, and interest) flowing from pass-throughs. We estimate investment income earned by each income group via pass-throughs as follows. First, we compute the fraction of taxable interest that, on aggregate, derives from pass-through businesses by dividing the total amount of interest distributed by pass-throughs to individuals by the total amount of taxable interest reported on individual income tax returns in 2011. We proceed similarly for dividends and capital gains. Second, we assume that these aggregate ratios are constant across the income distribution and were the same in 2007 as in 2011. Investment income flowing to individuals from partnerships in 2011 is from Cooper et al. (2016); investment income flowing from S-corporations is from SOI tabulations of S-corporation tax returns, see Saez and Zucman (2020).

earners this fraction increases to 35%. Concretely, this means that audits of associated pass-through entities may be required to comprehensively examine about 35% of the income reported by the highest-income taxpayers.

Third, the bias affects several income categories, not only partnership and S-corporation business income. As shown by Figure 7b, the income that flows to individuals from pass-throughs does not only consist of business income, but also includes dividends, interest, rents, and capital gains. Determining compliance for such pass-through investment income (reported on Schedule K-1) is more resource-intensive than for investment income reported (on 1099 Forms) by third parties directly to the individual and the IRS.[32]

3.2.2 Methodology to Estimate Evasion in Pass-Through Businesses

To assess the magnitude of the bias in the NRP due to undetected pass-through business evasion, we adopt the following strategy. We first make simple assumptions on the overall under-reporting rates for the various forms of pass-through income (business income, dividends, interest, capital gains). We consider a benchmark scenario and a range of alternative assumptions, both described below. We then assume that undetected pass-through income is distributed like reported pass-through income. For example, if the top 1% (by reported income) earns 60% of reported pass-through business income, we assume that the top 1% (by corrected income) also earns 60% of unreported pass-through business income. One could adopt a more sophisticated approach, especially for partnerships where the complexity of the partnership structure may be correlated with the rate of noncompliance, likely leading to more noncompliance at the top.[33] Last, we add pass-through evasion to the amount of under-reported income uncovered in the NRP by income bin, before DCE adjustment. To avoid any double counting with noncompliance detected in the NRP, we remove all pass-through-business-level evasion uncovered in the context of NRP audits.[34]

In our benchmark scenario, we assume that 20% of total pass-through business income, 5% of pass-through capital gains, and 3% of pass-through dividends and interest are under-reported. The rates assumed for pass-through investment income are the rates observed in the NRP for investment income directly earned by individuals (Appendix Table A1). For business income, our benchmark assumption is motivated by the following facts. First, as shown in Appendix Table A1, random audit estimates without DCE adjustment suggest that 23% of sole proprietorship income earned by the top 1% is under-reported

[32] A further complication is that income that should be taxed as, e.g., business income may be reported as capital gains to achieve the lower tax rate on long-term capital gains. There is potential for legal avoidance along these lines, noncompliance and, possibly, some gray area.

[33] Our methodology understates top-end evasion in two ways. First, if noncompliance is increasing in the complexity of the pass-through structure and such complexity increases with income, our procedure underestimates pass-through noncompliance among the wealthiest taxpayers. Second, our approach does not factor in any re-ranking from adding undetected pass-through business evasion to reported income.

[34] Specifically, 57.6% of partnership and S-corporation income evasion detected in the NRP is associated with an entity pick up (i.e., an audit of the corresponding business). Therefore we remove 57.6% of detected partnership and S-corporation evasion.

(37% for all sole proprietorship income). The under-reporting rate for C-corporation income estimated by the IRS (2019) is slightly lower, at 18%. As the size and sophistication of partnerships and S-corporations is between that of the sole proprietorships owned by top earners and that of C-corporations, assuming an under-reporting rate for pass-through business income in the 18%–23% range is reasonable. Second, the National Income and Product Accounts use a 20% income misreporting rate for C- and S-corporations combined. Third, the small-scale 2004 random audit study of S-corporations found an income misreporting percentage of 12% to 16% with no corrections for undetected evasion. Last, the 1982 TCMP estimated an income under-reporting rate of 26% for partnerships, albeit in a different tax system than today (GAO, 1995).

For our sensitivity tests, we allow the under-reporting rate of pass-through business income to vary from 12% (the lowest uncorrected rate found in the S-corporation random audit study) to 28% (slightly higher than the evasion rate for top 1% sole proprietorship income detected in the NRP before DCE). We allow the under-reporting rate of pass-through investment income to range from 0% to up two times the benchmark rate, i.e., up to 10% for capital gains and 6% for dividends and interest. We also implement additional robustness checks to account for traces of evasion specific to pass-through businesses. In one variant, we randomly disallow 20% of declared pass-through business losses, setting income to zero instead of the reported negative amount. This is motivated by the fact that a significant share of adjustments to partnership income in operational audits consists of disallowed losses, which thus appear to be correlated with noncompliance. In another variant, we classify part of the business income earned by circular partnerships—partnerships owned by other partnerships, which are owned by other partnerships, etc.—as tax evasion. Cooper et al. (2016) find that 15% of partnership business income is owned in circularly-owned partnerships in 2011; we consider the implications of classifying two-thirds of this income (i.e., 10% of all partnership business income) as tax evasion.

3.2.3 Implications of Pass-Through Evasion for the Distribution of Noncompliance

Figure 8a shows the results obtained with our benchmark assumptions. On aggregate, the pass-through correction (1.3% of true income) is almost twice as large as the offshore correction (0.7%). Accounting for pass-through businesses increases under-reported income at the top of the income distribution significantly. In uncorrected NRP data, the fraction of true income which is under-reported falls from 4% around the 90^{th} percentile of the (exam-corrected) income distribution to less than 1% in the top 0.01%. After adding pass-through under-reported income, the fraction of true income which is under-reported rises from 4% around the 90^{th} percentile to about 8% from the 99.5^{th} to the 99.95^{th} percentile. It then falls back to around 4% in the top 0.01%. This drop-off is due to the increasing prevalence of capital gains (as opposed to business income) at the very top; capital gains are assumed to have a low under-reporting rate (5%) in our benchmark

scenario. Tax evasion below the 90^{th} percentile is not significantly affected by the pass-through correction, because pass-through income is a negligible source of income for these groups.

To further assess the effect of accounting for pass-through evasion, it is useful to consider the top 1% as a whole. When including only the pass-through evasion uncovered in the NRP, the top 1% under-reports 2.6% of its true income, before DCE adjustment—a lower rate than the average under-reporting rate of 3.9%. By contrast, after accounting for pass-through evasion using our benchmark assumptions, the top 1% under-reports 6.7% of its income—a higher rate than the average rate of 5.0%. In other words, the pass-through correction alone removes the downward-sloping pattern of under-reporting rates found in uncorrected random audit data.

It is worth noting that factoring in undetected pass-through evasion also increases noncompliance at the top relative to DCE-corrected NRP data. Although the DCE adjustment may be seen as picking up part of this undetected unreported pass-through income, this adjustment appears quantitatively insufficient, especially at the very top end. As we have seen, detected pass-through business income evasion in the NRP is small (due to the practical limits faced by the audit process described above), and the DCE adjustment is proportional to detected evasion. As shown by Appendix Table A1, we estimate that only 2.0% of pass-through business income in the top 1% is estimated to be under-reported in the NRP, which is 10 times smaller than under-reporting in a reasonable benchmark scenario (20%). We note, however, that there is some potential overlap between the DCE adjustment and our proposed adjustment for pass-through evasion. We will take this overlap into account when we present our new income and tax gap estimates (combining the DCE-adjusted NRP and sophisticated evasion) in Section 4 below.

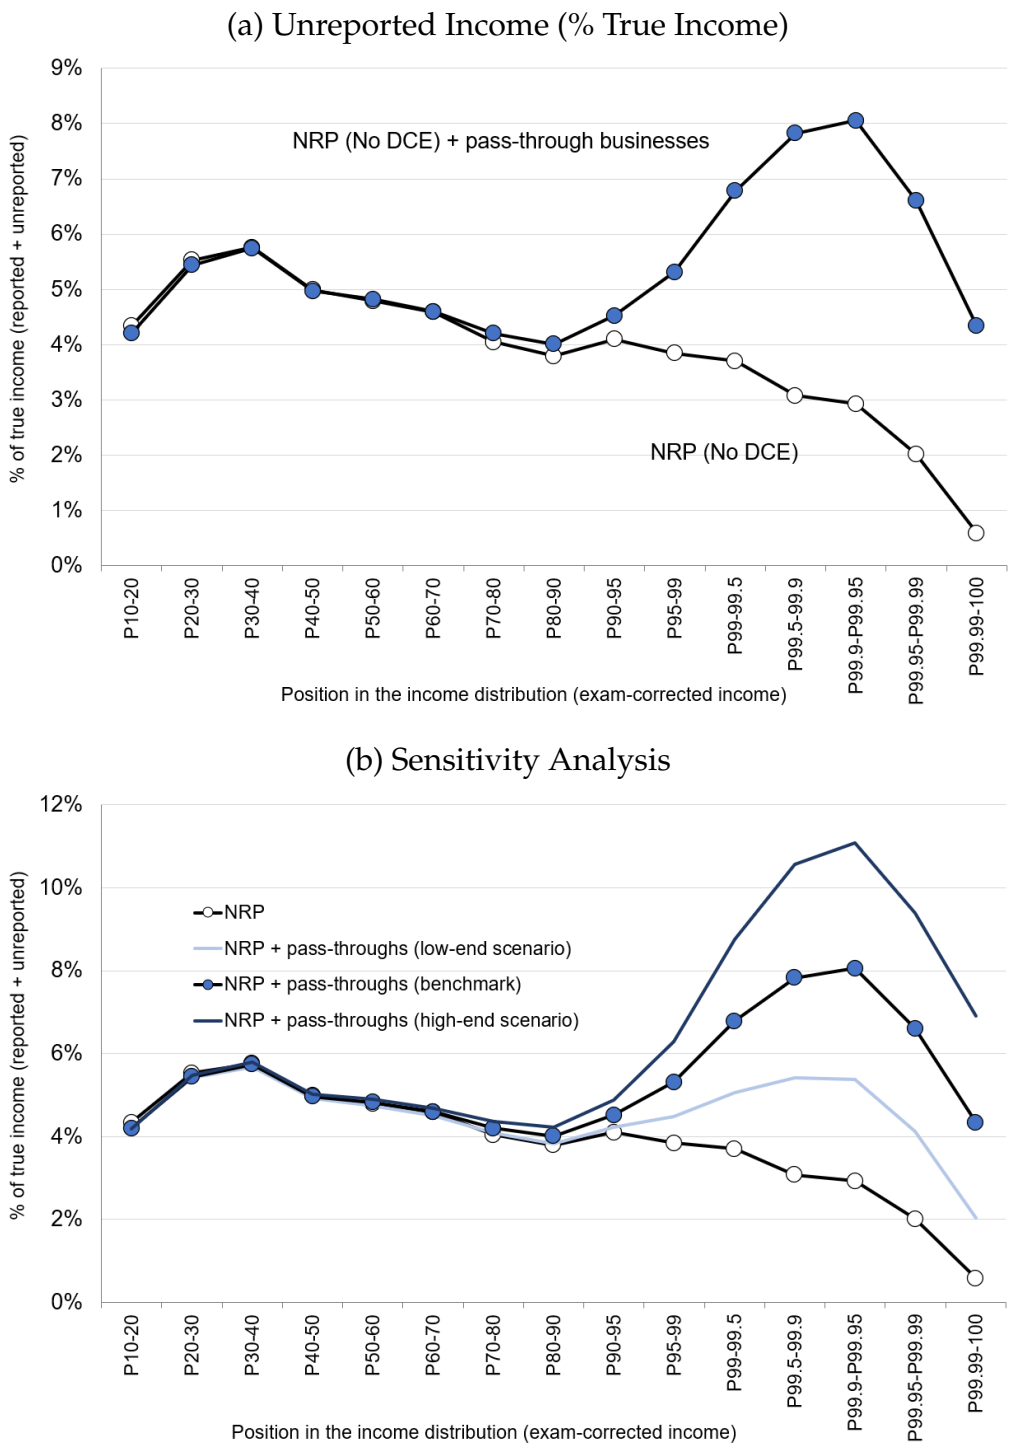
3.2.4 Pass-Through Evasion: Sensitivity Analysis

We now present the results from our sensitivity analysis. The main finding is that top-end pass-through evasion could be significantly higher than reported in our benchmark scenario, which should be seen as conservative.

First, as shown by Appendix Figure A12, accounting for pass-through business losses adds about one percentage point to the ratio of under-reported income to true income at the top. For the top 1% as a whole, the ratio of under-reported income to true income rises from 6.7% in our benchmark scenario to 7.5% when randomly disallowing 20% of reported pass-through losses. The effect of this correction is concentrated at the top because of re-ranking: taxpayers whose losses are disallowed typically are in the bottom decile of the reported income distribution, but end up in the top 1% of the corrected income distribution after their declared losses are set to zero.

Second, as shown by Appendix Figure A13, accounting for the income in circular partnerships adds 0.7

FIGURE 8: ACCOUNTING FOR PASS-THROUGH BUSINESS EVASION



Note: This figure shows estimates of unreported income by income group in the raw NRP (before DCE adjustment) and after adding estimates of pass-through business evasion. Taxpayers are ranked by exam-corrected income in NRP data, and pass-through adjustments are made on the basis of reported market income; this is the best available estimate of "true income" before DCE adjustments. In our benchmark scenario (top panel), we assume that 20% of pass-through business income, 5% of pass-through capital gains, and 3% of pass-through interest and dividends are under-reported, and that under-reported pass-through income is distributed like duly reported pass-through income. We remove all business-level pass-through evasion detected in the NRP before adding our estimates of business-level pass-through evasion. In the bottom panel, we report a high-end scenario in which 28% of pass-through business income, 10% of pass-through capital gains, and 6% of pass-through dividends and interest are unreported, and a low-end scenario in which only 12% of pass-through business income is unreported, while all pass-through investment income is duly declared.

percentage point to the ratio of under-reported income to true income in the top 1%. Partnership income is highly concentrated: 85% of partnership income went to the top 1% highest income-earners in 2007. The fact that 15% of partnership income appears to end up in circular partnerships (Cooper et al., 2016), of which we assume in Appendix Figure A13 that two-thirds is evasion, suggests that noncompliance rates at the top end may be significantly larger than in our benchmark estimates.

Third, our benchmark scenario assumes that 20% of pass-through business income is under-reported. This is somewhat conservative given the limited third-party reporting involved, especially for single-owner S-corporations and partnerships. Absent third-party-reported information, tax evasion rates tend to be high (IRS, 2019). Appendix Figure A14 investigates the effect of alternative assumptions on pass-through business income evasion. In a high-end scenario where 28% of pass-through business income is under-reported, under-reporting rises to 8.7% of true income for the top 1% as a whole. In a low-end scenario where 12% of pass-through business income is under-reported, the top 1% under-reports 5.0% of its true income, still more than average (4.5%).

Last, Appendix Figure A15 varies the assumed under-reporting rate of pass-through investment income. Our benchmark scenario assumes that 5% of pass-through capital gains are under-reported. There is reason to believe that this number is on the low-end. For some assets owned by pass-through businesses—such as real estate, business structures and equipment, and foreign securities—the IRS does not know the purchase price of the assets, making tax evasion possible. Before the implementation of the Foreign Account Tax Compliance Act, there was also little reporting on payments of foreign interest and dividends. It is thus informative to consider a variant where, starting from our benchmark scenario, the under-reporting rate of pass-through investment income is multiplied by two, keeping everything else the same. We find that the implied under-reporting rate rises by 1 percentage point for tax units in the top 0.01%. We also consider the opposite scenario in which all pass-through investment income is duly declared. As shown by Figure A15, accounting for pass-through business income evasion still makes a large difference at the top compared to the uncorrected NRP data.

Figure 8b summarizes our sensitivity analysis by depicting two scenarios: a high-end scenario in which 28% of pass-through business income, 10% of pass-through capital gains, and 6% of pass-through dividends and interest are unreported; and a low-end scenario in which only 12% of pass-through business income is unreported, while all pass-through investment income is duly declared. Even in the low-end scenario, income unreported by the top 1% is almost multiplied by two relative to what is detected in the NRP.

4 The Distribution of Noncompliance in the United States: Sophisticated vs. Less Sophisticated Evasion

In this section, we present new estimates of the distribution of noncompliance in the United States. These estimates adjust the methodology used by the IRS in its tax gap studies (e.g. IRS, 2016, 2019) to include the sophisticated forms of evasion that (given the tools, resources, and information constraints) can go undetected in random audits. We provide a benchmark scenario and extensive sensitivity analysis. The three main findings are the following. First, although our benchmark series feature only slightly more evasion on aggregate than in the standard IRS methodology, our proposed adjustments have large effects at the top of the income distribution. Specifically, our adjustments increase the aggregate tax gap by a factor of 1.1, but by a factor of 1.4 for the top 0.1% and by a factor of 1.7 for the top 0.01%. Second, in our benchmark series, we find clear evidence that under-reported income rises with income until at least the 99.95^{th} percentile of the true income distribution. Third, across all our specifications, we find that accounting for unreported income increases the top 1% income share significantly in 2006–2013.

4.1 Methodology for Allocating Unreported Income

In all our scenarios, we classify evasion in two categories: sophisticated and less sophisticated. In our benchmark scenario, we estimate less sophisticated evasion as in the standard IRS tax gap methodology and in Johns and Slemrod (2010). That is, less sophisticated evasion is the amount of noncompliance detected in NRP random audits, adjusted to account for differences in experience (and other observable characteristics) across examiners using the DCE methodology (Figure 3). Next, we estimate sophisticated evasion as the sum of offshore and pass-through evasion, using our benchmark scenario for each (Figure 6a and Figure 8a respectively). We then combine sophisticated and less sophisticated evasion in the simplest manner possible, i.e., assuming no re-ranking when adding them. We account for the potential overlap between the DCE adjustment and our proposed adjustment for sophisticated evasion. Specifically, we remove 57.6% of the DCE-adjusted estimate of partnership and S-corporation evasion in the NRP before adding our estimate of pass-through evasion.[35] In our benchmark scenario, the total amount of under-reported income is 14.6% on aggregate: 1.7% from sophisticated evasion and 12.9% from less sophisticated evasion.[36]

[35] As we have seen (footnote 34), 57.6% of partnership and S-corporation income evasion detected in the NRP is associated with an entity pick up (i.e., an audit of the corresponding business). Therefore, up to 57.6% of DCE-adjusted pass-through income evasion in the NRP can be seen as capturing business-level evasion. We remove all this business-level evasion before adding our own estimate of pass-through-business-level evasion.

[36] Total under-reported income (as a fraction of true income) for less sophisticated evasion (12.9%) is slightly lower than total DCE-adjusted NRP evasion (13.9%) because (i) we remove 57.6% of pass-through evasion; (ii) the "true income" denominator is enlarged by adding sophisticated evasion. Similarly, total under-reported income for sophisticated evasion (1.7%) is slightly lower than total offshore plus pass-through evasion (1.9% as a fraction of exam-corrected, non-DCE adjusted income) because the denominator is enlarged by the DCE adjustment.

This benchmark scenario is motivated by the following considerations. First, random audits are the most powerful tool available to detect the less sophisticated forms of evasion. Moreover, the fact that more experienced auditors systematically detect more evasion suggests that some evasion is missing in the raw exam-corrected data. The DCE methodology is the best methodology currently available to capture the insight that examiners vary in their capacity to uncover noncompliance; thus the DCE-corrected NRP is a natural starting point for allocating less sophisticated evasion. Last, as we have seen, due to the practical limits inherent to the conduct of random audits, the NRP largely misses offshore and pass-through business evasion. This is true even of the DCE-corrected NRP, because detected noncompliance on dividends, interest, capital gains, and pass-through business income is low, and the DCE correction is essentially proportional to detected evasion.

Because measuring evasion—especially undetected evasion—necessarily involves a margin of error, we also consider a large number of robustness tests capturing the key dimensions of uncertainty. First, we implement the high-end and low-end scenarios for offshore and pass-through evasion described in Figures 6b and Figure 8b, respectively. Second, we consider the case where sophisticated evasion on aggregate is lower or higher than 1.7% of true income but distributed like in our benchmark series. At the low-end, we assume that sophisticated evasion adds up to 1.2% of true income. This reflects a post-FATCA world in which FATCA would be fully effective (so that under-reporting rates on offshore capital income would be reduced by a factor of 4 relative to our benchmark scenario and become similar to the under-reporting rates on onshore capital income) while other forms of sophisticated evasion would remain constant. At the high-end, we consider a scenario in which total sophisticated evasion adds up to 2.2% of true income. This scenario reflects the hypothesis that there may be quantitatively significant forms of sophisticated evasion other than pass-through and offshore evasion, such as abusive uses of trusts, charities, and tax shelters used by high-net-worth individuals. Last, in keeping with our approach so far, we also show a version of all our results without the DCE adjustment. As we shall see, the finding that accounting for unreported income increases the top 1% income share is robust to all these variations.

4.2 Main Results

We start by describing the income under-reporting results from our benchmark scenario, which are depicted in Figure 9. In our preferred estimate (blue line), we find that income under-reporting rises sharply with income until roughly the 99.95^{th} percentile of the distribution of true income. Specifically, unreported income is around 7% of true income in the bottom half of the distribution, rises slowly to close to 10% for the middle-class (from the median to the 90^{th} percentile), and then rises sharply after the 95^{th} percentile.

It stabilizes at 25% in most of the top 1%, and then declines to 15%–20% in the top two income bins we consider (P99.95 to P99.99 and top 0.01%).

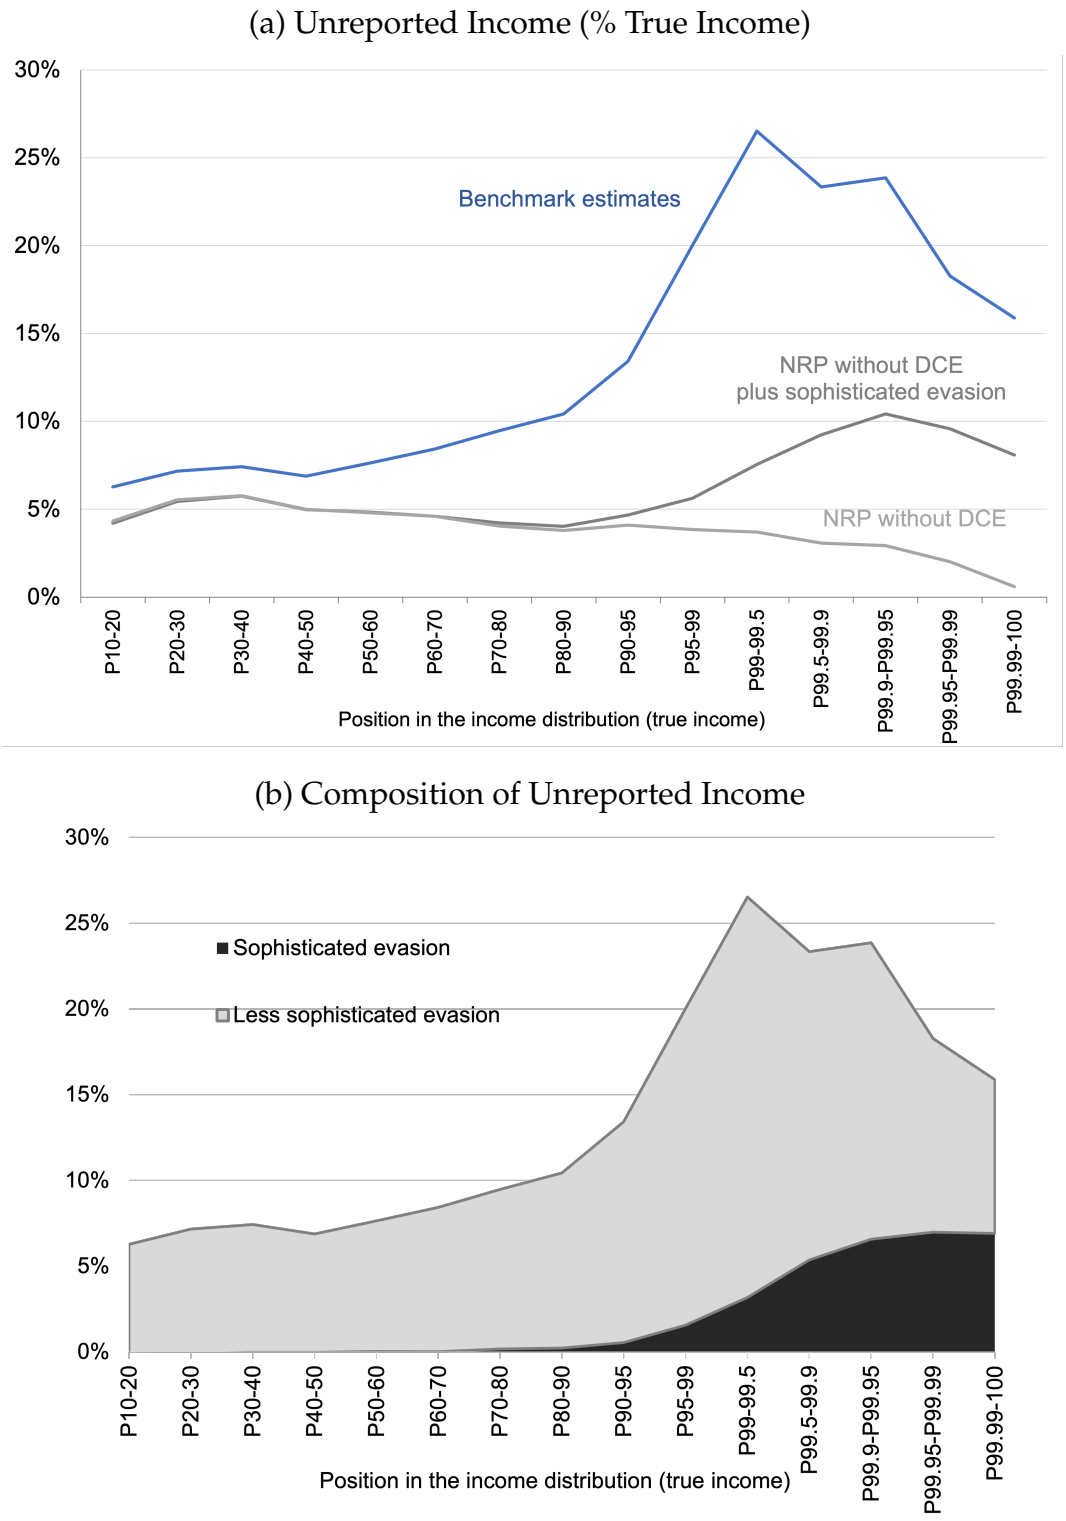
Figure 9b decomposes our benchmark estimate of evasion into sophisticated versus less sophisticated evasion. As one moves up the income distribution, sophisticated evasion rises while less sophisticated evasion converges to zero. More precisely, sophisticated evasion rises from 0% in the bottom 90% to 8% of true income in the top 0.01%. Meanwhile unsophisticated evasion first rises with income (as opportunities for such evasion increase, e.g., on sole proprietorship and rental income) and then falls, as income that can be easily concealed becomes less important.

A number of additional results are worth noting. First, we compare our results to those obtained following the standard IRS methodology (e.g. IRS, 2019; Johns and Slemrod, 2010). On aggregate, our benchmark income under-reporting gap is only about 1.1 times larger than the gap found using the DCE-adjusted NRP (see Appendix Table A5). However, our benchmark adjustment increases the income under-reporting gap by a factor of 1.4 for the top 0.1%. This is because sophisticated evasion, although not large on aggregate in our benchmark estimates, is highly concentrated, with 40% of it going to the top 0.1%. Second, in the spirit of the analysis presented in Section 3, in Figure 9a we report how sophisticated evasion modifies the raw (i.e., non-DCE corrected) profile of evasion detected in the NRP. A striking finding emerges. While in raw NRP data, the highest earners under-report less income (relative to true income) than other taxpayers, after adding sophisticated evasion the opposite is true. Under-reported income rises from 4%–5% of true income in the bottom 95% of the income distribution to 8%–10% in the top 0.5%.

Table 2 reports the share of total unreported income attributable to various income groups. In our benchmark scenario, 33.9% of all unreported income is earned by the top 1% of the true income distribution. When disregarding sophisticated evasion (i.e., in the DCE-adjusted random audit data, as used, e.g., by Johns and Slemrod (2010)), unreported income is slightly less concentrated: the top 1% earns 30.3% of it on average over 2006–2013. This is slightly higher than the comparable 27% figure found in the 2001 NRP by Johns and Slemrod (2010). We also report the distribution of unreported income when discarding the DCE adjustment altogether. In that case, without sophisticated evasion, the top 1% earns 12.8% of all under-reported income; adding sophisticated evasion, the top 1% earns 32.5% of all unreported income, close to the standard DCE-adjusted NRP number of 30.3%. In other words, as far as the distribution of noncompliance is concerned, modifying the raw evasion detected in random audits by using the DCE methodology and modifying the raw evasion detected in random audits by adding sophisticated evasion deliver similar results. In both cases, the top 1% earns close to a third of all unreported income.

Appendix Figure A16 presents our benchmark tax gap estimates. We find that taxes evaded rise from about 20% of taxes owed in the bottom 95% of the distribution to 25%–30% in most of the top 1%, before

FIGURE 9: THE DISTRIBUTION OF NONCOMPLIANCE IN THE U.S.: BENCHMARK ESTIMATES



Note: This figure shows estimates of under-reported income by true income groups, when combining DCE-adjusted NRP evasion and our benchmark estimate of sophisticated evasion (offshore and pass-through business evasion) in 2007. For comparison, the top panel reports raw (i.e., before DCE adjustment) evasion detected in the NRP, and raw evasion combined with sophisticated evasion. We rank individuals by estimated true income either before or after DCE adjustment. For details on raw evasion detected in the NRP, NRP evasion after DCE adjustment, benchmark offshore evasion, and benchmark pass-through evasion, see notes to Figure 2, Figure 3, Figure 6 and Figure 8 respectively.

TABLE 2: SHARES OF UNREPORTED INCOME, 2006-2013, IN % OF TOTAL UNREPORTED INCOME

	NRP No DCE No sophisticated	NRP With DCE No sophisticated	NRP No DCE Add sophisticated	Our benchmark
P0-10	1.4	0.8	1.0	0.7
P10-20	1.3	0.5	0.8	0.4
P20-30	2.7	0.9	1.8	0.8
P30-40	4.1	1.3	2.8	1.2
P40-50	5.2	1.8	3.5	1.7
P50-60	6.9	2.7	4.7	2.6
P60-70	9.4	4.4	6.5	4.1
P70-80	11.8	7.2	8.5	6.8
P80-90	16.0	12.0	11.7	11.3
P90-95	12.3	11.5	9.7	11.0
P95-99	16.1	26.5	16.4	25.5
P99-99.5	3.9	9.6	5.7	9.2
P99.5-P99.9	5.3	11.5	11.6	12.5
P99.9-99.95	1.4	3.2	3.8	3.6
P99.95-P99.99	1.6	3.3	5.7	4.3
P99.99-100	0.6	2.8	5.7	4.3
Top 1%	12.8	30.3	32.5	33.9

Note: This table reports the distribution of unreported income across income groups, for different measures of unreported income. Tax units are ranked by their estimated true income (equal to reported income plus estimated unreported income). The first column shows the distribution of unreported income detected in the NRP with no adjustment for undetected evasion (in particular, without the DCE methodology). The second column shows the distribution of unreported income in the NRP after the DCE adjustment. The third column shows the distribution of unreported income detected in the NRP without DCE adjustment, but adds our benchmark estimate of sophisticated evasion. The last column shows the distribution of unreported income in our benchmark scenario which (as described in Section 4.1) incorporates both sophisticated evasion and the DCE adjustment.

falling back to close to 20% in the top 0.01%. Taxes evaded in the top 0.01% are 1.7 times as large as in DCE-adjusted random audits. That is, even the DCE methodology underestimates the true tax gap at the top by a factor of almost 2, according to our benchmark estimates. In our benchmark scenario, as shown by Appendix Table A7, we estimate that 41.8% of all federal income taxes unpaid are attributable to the top 1% (ranked by corrected income). For comparison, the top 1% (by reported income) paid 35.5% of all federal income taxes on average in 2006–2013. Fully collecting the unpaid income taxes of the top 1% would increase income tax revenue by an amount equivalent to 11.8% of the aggregate amount actually collected. For example, it would have increased tax revenue by $202 billion in 2019 (a year when actual collection was $1,713 billion, as reported in the National Income and Product Accounts), according to our benchmark estimates.[37]

The tax gap we estimate for the top 0.01% (close to 20% of taxes owed) is similar to the one obtained by

[37] For the years in our sample, we estimate that closing the tax gap for the top 1% would raise about $146 billion annually in 2012 dollars (see Tables A5 and A7). About $15 billion of the difference between this and the $202 billion figure for 2019 is due to inflation, while the rest is due to growth of top incomes.

Alstadsaeter et al. (2019) in Scandinavia by combining random audits with estimates of unreported offshore income (25% for the top 0.01% by wealth). However, the Scandinavian tax gaps estimated by Alstadsaeter et al. (2019) in the bottom 90% are much lower than in the U.S., around 5% through most of the distribution vs. 15%–20% in the United States. This is due to an number of reasons. First, the United States is unique in using the DCE methodology to adjust detected noncompliance in random audits; in other countries with random audit programs, researchers and tax authorities use uncorrected random audit data similar to those presented in Figure 2 above. Second, as already noted (footnote 8), taxes evaded as a fraction of taxes owed are mechanically high at the bottom of the distribution in the United States when (as done in Figure A16) refundable tax credits are subtracted from taxes owed. Removing refundable tax credits would reduce the tax gap at the bottom by about 7 percentage points. Third, in contrast to Scandinavian countries, many large businesses are organized as pass-throughs entities in the United States. Because pass-through business income is subject to the federal income tax, our estimates of the federal income tax gap include evasion by those businesses. To maximize international comparability, it would be desirable to estimate comprehensive distributional tax gaps covering all taxes, including the corporate tax—in effect treating all businesses, in both the United States and Scandinavia, as if they were pass-through businesses. We leave this important task to future research.

Finally, we present the results of our sensitivity analysis in Appendix Figure A17. The basic pattern of under-reporting depicted in Figure 9b remains across all the scenarios we consider. Under-reporting rises with income up to the P99–P99.5 group, then stabilizes at a high level (about 25%–30% of true income) in the P99.9–P99.95 group. In our upper-bound scenario for sophisticated evasion, evasion remains at that level up to the very top of the distribution; the top 0.01% under-reports about 25% of its true income. In the other scenarios, there is a drop-off in evasion in the top 0.05%.

4.3 Implications for the Measurement of Inequality

There has been renewed interest in recent decades in the study of income and wealth distributions. Income tax returns are a key data source for the study of inequality. However, the tax data used to study inequality in the United States are before any correction for unreported income. How does accounting for unreported income affect what we know about inequality?

Table 3 reports the distribution of taxable market income on average over 2006–2013 with different treatments of unreported income. The first column shows the distribution of reported income, that is, with no adjustment for evasion. The second column shows the distribution of reported plus detected underreported income (i.e., with no adjustment for undetected evasion). This distribution is very similar to the

distribution of reported income, because only 3.9% of income is found to be under-reported on aggregate in the raw NRP. The third column shows the distribution of reported plus DCE-adjusted under-reported income. This is the same statistic as the one reported for the year 2001 in the first column of Table 5 in Johns and Slemrod (2010).[38] We observe a notable shift in income concentration. While the top 1% (by reported income) earns 20.3% of reported income, the top 1% (by exam-corrected, DCE-adjusted income) earns 22.2% of true income. Thus accounting for tax evasion increases income concentration substantially in 2006–2013. Finally, the last column of Table 3 shows the distribution of true income in our benchmark scenario. We observe a further shift in income concentration: the top 1% earns 22.8% of true income, 2.5 points more than when disregarding noncompliance.

The result that accounting for noncompliance increases the top 1% income share is robust to alternative assumptions about the size and distribution of undetected evasion from unadjusted random audit data. First, in the extreme case where there is zero sophisticated evasion, we are back to the DCE-adjusted estimates reported in column 3 of Table 3—i.e., the top 1% income share is 1.9 percentage point higher after accounting for noncompliance than before. In the case where there is zero offshore evasion (only pass-through business income evasion), the top 1% income share rises by more than 1.9 point. Second, if we discard the DCE methodology entirely and simply add sophisticated evasion to the raw evasion uncovered in the NRP, the top 1% income share rises from 20.1% pre-audit to 21.1% (col. 4 of Table 3). The increase is more muted than in the standard IRS tax gap methodology, because there is much less under-reported income on aggregate in that scenario (5.7% of true income vs. 13.9% in the DCE-adjusted NRP).

These results have implications for estimating the distribution of total U.S. national income, as in Piketty et al. (2018) and Auten and Splinter (2019). For the computation of U.S. national income, the Bureau of Economic Analysis includes income which should be reported in tax returns but is not. For wages, sole proprietorship income, and partnership business income, the amounts of unreported income included in national income are directly based on the DCE-adjusted NRP.[39] For other income categories, the inclusion of unreported income is implicit; for instance, aggregate rental income is higher in the National Income and Product Accounts (NIPAs) than in tax returns in part because some rental income is unreported in individual income tax returns.[40] Because the top 1% income share is higher after allocating unreported income in the DCE-adjusted NRP (and even more so after accounting for sophisticated evasion), our results in this paper suggest that a proper treatment of tax evasion increases income concentration in 2006–2013.

[38] See Appendix Table A4 for a comparison of our results with Johns and Slemrod (2010).

[39] In 2013, the adjustment amounts to $76.8 billion for wages (NIPA Table 7.18 line 2) and $639.8 billion for sole proprietorship and partnership income (NIPA Table 7.14 line 2).

[40] No explicit reconciliation between aggregate rental income in the NIPAs and in 1040s is provided by the Bureau of Economic Analysis because NIPA rental income is not estimated based on tax returns. Other conceptual differences (e.g., the treatment of depreciation) contribute to the gap between NIPA income and income reported in tax returns; see Saez and Zucman (2020, section 3.1.1).

TABLE 3: SHARES OF TRUE INCOME, 2006-2013, IN % OF TOTAL INCOME

	NRP Before exam	NRP After exam No DCE No sophisticated	NRP After exam With DCE No sophisticated	NRP After exam No DCE Add sophisticated	Our benchmark
P0-10	-2.3	-1.9	-1.5	-1.8	-1.4
P10-20	1.3	1.2	1.0	1.2	1.0
P20-30	2.2	2.0	1.7	1.9	1.7
P30-40	3.0	2.9	2.5	2.8	2.4
P40-50	4.3	4.2	3.6	4.1	3.6
P50-60	5.9	5.8	5.0	5.7	5.0
P60-70	8.4	8.2	7.3	8.1	7.2
P70-80	11.8	11.7	10.8	11.5	10.7
P80-90	16.8	16.9	16.2	16.7	16.0
P90-95	11.9	12.1	12.2	11.9	12.1
P95-99	16.5	16.8	19.0	16.8	18.9
P99-99.5	4.2	4.3	5.2	4.3	5.2
P99.5-P99.9	6.9	6.9	7.8	7.2	7.9
P99.9-99.95	2.0	2.0	2.2	2.1	2.3
P99.95-P99.99	3.3	3.2	3.3	3.4	3.4
P99.99-100	4.0	3.8	3.8	4.1	4.0
Top 1%	20.3	20.1	22.2	21.1	22.8

Note: This table reports the distribution of true market income (defined as total income reported on form 1040 minus Social Security benefits, unemployment insurance benefits, alimony, and state refunds) for different measures of unreported income. Tax units are ranked by their estimated true income (equal to reported income plus estimated under-reported income). The first column shows the distribution of reported income (i.e., with no adjustment for evasion). In the second column, under-reported income only includes what is detected in raw random audits with no adjustment of any kind (in particular no DCE adjustment). In the third column, under-reported income is equal to DCE-adjusted NRP unreported income. In the fourth column, under-reported income is equal to raw NRP unreported income (with no DCE adjustment) plus sophisticated evasion. The last column shows the distribution of true income in our benchmark scenario combining the DCE adjustment and sophisticated evasion.

In Auten and Splinter (2019), the top 1% income share is lower by 0.8 percentage point on average over 2006–2013 after the allocation of unreported income.[41] In Piketty et al. (2018) the top 1% income share is higher by 0.7 percentage points after including the forms of evasion explicitly identified in the NIPAs.

[41] For example, the top 1% income share before the allocation of under-reported income is 16.9% in 2013 (Auten and Splinter, 2019, Table C1-Incomes, col. DE divided by col. DB) and 16.1% after adding unreported income (Auten and Splinter, 2019, Table C1-Incomes, col. DM divided by col. DJ). We identify two drivers of this drop: the allocation of under-reported income among filers and the treatment of non-filers. First, Auten and Splinter (2019) allocate unreported income based on the distribution of unreported income by bins of reported income in the 2001 NRP (Johns and Slemrod, 2010). A large fraction of unreported income is earned by tax units who appear to be in the bottom 99% of the reported income distribution (which is not surprising, since noncompliance reduces their reported income). Allocating noncompliance based on its distribution by bin of reported income thus boosts the bottom 99% income share. In reality, if a tax unit reports, e.g., $0 in income and has $1 million in unreported income, this tax unit is in the top 1% of the true income distribution and $1 million in unreported income should be allocated to the top 1%, not to the bottom 99%. Auten and Splinter (2019) incorporate some re-ranking in their methodology to account for this, but their procedure falls short of capturing the full extent of the re-ranking observed in the NRP, especially in recent years. Specifically, their methodology results in a small decline in the top 1% income share in 2001 and a larger decline in more recent years (as opposed to no change in the 2001 NRP studied by Johns and Slemrod (2010) and an increase in 2006–2013 in our data). That is, their methodology generates a decrease in top income shares over time. Second, Auten and Splinter (2019) allocate 15% of their total under-reported income aggregate to non-filers, with no re-ranking; i.e., all this income is allocated to the bottom of the distribution. This also causes a drop in the top 1% share. We do not analyze non-filer income in our paper, but we note that recent analysis shows that high-income non-filers drive the bulk of the non-filer tax gap in recent years (TIGTA, 2020).

We note that the data in Tables 2, A5, A6, and A7 can be used to adjust for income under-reporting and tax due in distributional macroeconomic statistics even when without data on corrected incomes. Doing so requires carefully accounting for re-ranking, as reported incomes are negatively selected on noncompliance. Methods for accounting for re-ranking based on NRP data, not including the sophisticated evasion we study here, are being developed in work in progress by Auten and Langetieg (2020). Our adjustments for sophisticated evasion specifically are straightforward to add on top of any existing method for adjusting the distribution of income to account for mis-reported income, because our sophisticated adjustments are done on the basis of reported incomes. In other words, one could simply add our estimated amount of sophisticated evasion in each (reported) income bin, based for instance on Table A6, after implementing existing methods for distributing other under-reported income.[42]

Accounting for tax evasion also affects the measurement of wealth inequality. Following the work of Saez and Zucman (2016), a number of authors have estimated U.S. wealth inequality by capitalizing income tax returns. Accounting for sophisticated tax evasion can increase estimated top wealth shares in at least two ways. First, a large fraction of offshore wealth (foreign bank accounts, portfolios of equities and bonds held through foreign financial institutions, holdings of foreign mutual funds, foreign real estate) is not captured in the official Federal Reserve estimates of the total amount of U.S. household wealth (Zucman, 2013). As offshore wealth is highly concentrated, accounting for it increases top wealth shares; for instance in the supplementary series of Saez and Zucman (2016) it increases the top 0.1% wealth share by 0.9 percentage point in 2012 (see Saez and Zucman, 2016, p. 539 for a discussion). Second, accounting for pass-through business evasion can also increase top wealth shares. In Smith et al. (2019), the aggregate market value of pass-through business wealth and its distribution are obtained by capitalizing reported business income (gross of interest paid and depreciation), sales, and assets. If pass-through businesses under-report income to their owners (e.g., by understating sales, or by retaining income in circular partnerships), aggregate business wealth is downward-biased. Given that business wealth is highly concentrated, factoring in business tax evasion would increase wealth concentration.[43]

[42] When adding sophisticated evasion to estimates of under-reporting without DCE corrections, we recommend using the "Sophisticated – after exam" column of Table A6. When adding sophisticated evasion to DCE-adjusted estimates we recommend the "Our benchmark – DCE-corrected" column. The difference between the two derives from our modification of DCE-adjusted pass-through business under-reporting to avoid double counting; see also the notes to Tables A5 and A6. We also note that care should be taken in applying these adjustments in different periods. Outside our sample period of 2006-2013, our view is that a multiplicative adjustment for sophisticated evasion rather than an additive one is preferable, to remain neutral about the dynamic evolution of under-reporting-adjusted top income shares. A multiplicative adjustment ensures that if reported income shares do not change over time, neither do under-reporting-corrected income shares. We have emphasized the multiplicative figures elsewhere in the main text for this reason.

[43] In the Saez and Zucman (2016) methodology, the aggregate market value of pass-through businesses is taken from the official Federal Reserve Financial Accounts. Aggregate wealth in these accounts is not estimated based on reported business income, so in principle it is not affected by business income tax evasion (although it might be too high or too low for other reasons). If unreported pass-through income is proportional to reported pass-through income (as our benchmark estimates assume), then the wealth distribution in the Saez and Zucman (2016) methodology is not affected by business income tax evasion.

5 Theory

In this section, we seek to inform the interpretation of the above body of empirical results with some simple economic theory. Our central goal is to explain why some forms of evasion, like offshore evasion and pass-through evasion are 1) concentrated at the top of the income distribution, and 2) difficult to detect by random audit. We also briefly consider some related optimal tax enforcement considerations.

5.1 Individual Evasion Decisions

Setup. As in Allingham and Sandmo (1972) an individual determines how much income to report to the tax authority out of exogenous true income, y. Evaded income is denoted by e. The only modification is that we give the agent the option to take a binary concealment action $a \in \{0, 1\}$ that will reduce the probability of detection of evasion at some fixed cost κ. The individual's optimization problem is:

$$\max_{e \in [0,y], a \in \{0,1\}} \quad (1 - p(a))u((1-\tau)y + \tau e - \kappa a) + p(a)u((1-\tau)y - \tau\theta e - \kappa a), \tag{1}$$
$$\text{subject to } (1-\tau)y - \tau\theta e - \kappa a > 0$$

where τ is the tax rate, $\theta > 0$ is the penalty rate (following Yitzhaki, 1974), $u()$ is standard risk-averse utility over final consumption with $u' > 0, u'' < 0$, and $p(a)$ is the probability of detection with $p(1) < p(0)$. The given constraint requires that final consumption is positive in both states.[44] If we restrict the individual to $a = 0$, this model obviously becomes the Allingham-Sandmo model.

Because we view this as an intuition-building model, there are multiple interpretations of the action a. For example, the action one has in mind could be shifting income offshore or adopting a complex business structure with some income allocated to offshore or tax-exempt entities. With this interpretation, one can think of $p(a)$ as the probability of audit times the probability that the auditor discovers the taxpayers' evasion conditional on a, $p(a) = \Pr(\text{audit}) * \Pr(\text{auditor detects evasion}|a)$. Prior literature implicitly assumed the second term in this expression is one, even when adopting an endogenous detection probability (Kleven et al., 2011). We could also interpret a as the adoption of a gray area avoidance/evasion position that is uncertain to be successful if the position is legally challenged. In this case, the model looks exactly the same, but we would impose $p(a) = \Pr(\text{audit}) \Pr(\text{taxpayer loses legal challenge}|a)$.

For simplicity, we do not directly model third-party information, though it has been shown to be important and straightforward to incorporate into the Allingham-Sandmo framework (Kleven et al., 2011). Intuitively, when information in a taxpayer's income is acquired by the tax authority from a third party,

[44] We could add the constraint that $(1-\tau)y - \tau e - \kappa a > 0$ as well, but given that consumption is always weakly larger in the state where evasion is not detected, this constraint will never bind.

evasion on this income becomes easily detectable, and not reporting such income would have a very high probability of detection. One can essentially think of third-party information as imposing an upper bound on evasion, which would be a straightforward extension of the model here—note that we already require that $e < y$. One can also think of the shifting of income away from what would be covered by third-party information (for example shifting assets offshore to avoid third-party reporting on capital income by domestic financial institutions, before FATCA) as a concealment action.

Finally, we introduce some notation to facilitate exposition of the results. First, we denote the fraction of true income evaded—the analogue to the tax gap in the model—by $g = e/y$. To distinguish the chosen g in the optimization problem from Equation (1) from the g chosen under $a = 1$ or $a = 0$, we let $g_a(y, p)$ denote the level of g the taxpayer would choose if we maximize the objective restricting to $a = 0$ or $a = 1$, given true income y and probability of detection p. Second, Allingham and Sandmo showed that the effect of changes in income on evasion depend on absolute and relative risk aversion, which we denote by $A(c) \equiv -u''(c)/u'(c)$ and $R(c) \equiv -cu''(c)/u'(c)$, respectively.

Concealment at High Incomes. We begin by developing some intuition for how the tax gap should vary by income, and how this depends on whether the individual takes the concealment action. The following Lemma summarizes everything we know about this from Allingham and Sandmo's analysis of their original model, which recall is nested under $a = 0$.

Lemma 1. ***The Allingham-Sandmo Tax Gap.***

L1.1. *If the individual is risk averse, g_0 is decreasing over p, $\partial g_0/\partial p < 0$.*

L1.2. *If absolute risk aversion is decreasing, $A' < 0$, evasion e is increasing in true income y.*

L1.3. *If relative risk aversion is constant $R' = 0$, g_0 is constant over income, $\partial g_0/\partial y = 0$.*

L1.4. *If relative risk aversion is decreasing $R' < 0$, g_0 is decreasing over income, $\partial g_0/\partial y < 0$.*

L1.5. *If relative risk aversion is increasing, $R' > 0$, g_0 is increasing over income, $\partial g_0/\partial y > 0$.*

Proof. See Allingham and Sandmo (1972). □

Our goal is to build a model that captures the intuition that, as income grows very large, the fixed cost of adoption κ becomes a trivial share of income, so the taxpayer will opt for the lower detection probability given the trivial cost. The following assumption ensures that as the cost becomes a trivial share of income, the benefits of adoption do not also become trivial:

Assumption 1. *As y becomes arbitrarily large, $g_0(y, p_1)$ approaches a strictly positive constant.*

Assumption 1 is stated as an assumption about optimal behavior, but it actually imposes restrictions on the primitives of the model, especially risk preferences. It first requires that the limit of g_0 as y tends to infinity exists, which rules out some extremely strange risk preferences and behaviors (e.g. oscillations). Additionally, from Lemma 1, we know that Assumption 1 is satisfied under constant and increasing relative risk aversion, as these imply that g_0 is constant or increasing with income.[45] The main situation in which this assumption could fail is if relative risk aversion is decreasing at large incomes. However, decreasing relative risk aversion alone does not ensure violation of the assumption: g_0 could decrease but approach some strictly positive constant as y becomes large. Assumption 1 is unambiguously violated under constant absolute risk aversion, however: in this case e is constant over y and $g = e/y$ will become trivial as y grows large.

In any case, our empirical results above suggest that this assumption is realistic, at least for a sizable fraction of very high-income individuals, because concealment using the types of technologies we have in mind for $a = 1$ is widespread at the top of the distribution. Assumption 1 thus allows us to construct a theoretical argument that mirrors our intuition and empirical observation.

Lemma 2. *Under Assumption 1, as y becomes arbitrarily large, $g_1(y, p_1) - g_0(y, p_1)$ converges to zero.*

The proof of Lemma 2 and all subsequent results are in Appendix B. We solve the optimization problem (1) by first determining the optimal level of evasion under $a = 0$ and $a = 1$, and then comparing welfare under the two of them to decide whether to adopt. Lemma 2 shows that holding the probability of detection fixed at p_1, the fixed cost becomes irrelevant for behavior as y becomes large. This result helps us compare $a = 1$ and $a = 0$ for large y to determine which action the individual chooses. This comparison leads to our first main result.

Proposition 1. ***High-Income Concealment.*** *Under Assumption 1, there is a cutoff in the model $\hat{y}$ such that holding all else fixed, $y > \hat{y} \implies a = 1$ is optimal.*

The full proof of Proposition 1 is involved, so we sketch the intuition here. Adoption involves a trade-off between a lower probability of detection and the fixed cost. The fixed cost becomes trivial as a share of income y at large incomes. Lemma 2 states that because of this, behavior if the individual adopts the concealment technology is essentially unaffected by the fixed cost. At large incomes therefore, adoption incurs a trivial cost, but, by Assumption 1, the benefits of a lower probability of detection are non-trivial. The individual therefore adopts at sufficiently high income. Moreover, this logic applies even in the case

[45] This statement ignores the corner solution in which the individual never evades under $p = p_1$, for any y, so $g = 0$ everywhere. With respect to the results below, we note that in this case the individual would obviously never adopt $a = 1$. Henceforth we continue to ignore this corner solution.

where marginal utility u' becomes trivial for large y; covering this case is the main reason the proof is more involved than one might naively expect.

We note that Proposition 1 is not a unique cutoff rule, wherein individuals adopt *if and only if* true income exceeds some threshold. A wide parameter search of simulations of the model suggest that under constant relative risk aversion, optimal concealment does in fact follow a unique cutoff rule over y. However, we have not explicitly characterized the conditions under which such a single cutoff rule result obtains. In any case, it may be unrealistic to expect that concealment actions are always exclusively concentrated at the very top of the income distribution. For example, the use of cash to conceal transactions, potentially even from auditors, is generally believed to be widespread for self-employment income throughout the income distribution - see e.g. Slemrod et al. (2017). The result in Proposition 1 does however provide an explanation why in many cases, the adoption of complex and dubious sheltering strategies is concentrated at the very top.

Audits and Concealment. We now show that changing audit probabilities can induce the adoption of concealment. We do not condition this analysis on income for simplicity, but we discuss the results in relation to large empirical changes in audit probabilities for high-income taxpayers specifically.[46]

Proposition 2. ***Incentivizing Concealment.*** *Suppose a policy increases the probability of detection only if $a = 0$. This policy will increase concealment.*

Proposition 2 implies that if there is a concealment action that shields evasion against a particular type of enforcement, increasing that type of enforcement incentivizes adoption of that concealment strategy. Our results suggest that broad 1040 audits like NRP random audits in the U.S. tax system do not detect a number of different types of evasion, such as offshore evasion. Increasing these types of audits will therefore incentivize adoption of those types of evasion. It also implies that increasing more sophisticated types of audits will incentivize *even more sophisticated* types of concealment, if available. Additionally,the proposition implies that frequent audits can incentivize the adoption of gray-area avoidance strategies that would require protracted litigation to challenge. Altogether, the fact that audits overall are relatively common at the top of the income distribution (see Section 2.2) suggests that a variety of more sophisticated concealment and dubious avoidance activities should be more prevalent at the top, all else equal.

Implications for Tax Gap Estimation So far, we have given two reasons why the adoption of concealment strategies are especially common at the top of the income distribution. We next formalize the implications of this possibility for the estimation of the tax gap.

[46]Note that empirical audit probabilities are based on reported income not true income.

For simplicity, we assume that income is the only source of heterogeneity in our model; the basic point we make here does not rest on this assumption. Let $p_{D|a}$ denote the probability that a random audit detects evasion for $a \in \{0, 1\}$. Consistent with the model above, we suppose detection is less likely when $a = 1$: $p_{D|a=1} < p_{D|a=0}$. If concealment is optimal at income y, so $a(y) = 1$, then evasion is detected in a random audit with probability $p_{D|a=1}$, in which case the level of evasion the individual chooses under $a = 1$ is added to the tax gap. The *estimated tax gap* conditional on true income y is therefore $\hat{g}(y) = a(y)p_{D|a=1}g(p_1, \kappa, y) + (1 - a(y))p_{D|a=0}g(p_0, 0, y)$. Writing true evasion as $g = a(y)g(p_1, \kappa, y) + (1 - a(y))g(p_0, 0, y)$, we have that the bias in the estimated tax gap is

$$\hat{g}(y) - g(y) = a(y)(p_{D1} - 1)g(p_1, \kappa, y) + (1 - a(y))(p_{D0} - 1)g(p_0, 0, y). \tag{2}$$

The estimated tax gap is underestimated whenever random audits do not detect all evasion. More importantly, because $p_{D1} < p_{D0}$, the bias is larger when $a(y) = 1$. As our results above suggest that this is particularly likely at large values of y, equation (2) suggests that this bias is particularly large at the top of the distribution. Because $g(p_1, \kappa, y) > g(p_0, 0, y)$, it is even possible that the measured tax gap is decreasing in income while the true tax gap is increasing in income.

5.2 Implications for Tax Administration

In this section, we consider the problem of sophisticated tax evasion from the perspective of the tax authority. Our goal is to understand how the tax authority responds to the adoption of concealment strategies by certain taxpayers, which we model as an increase of the cost of collecting revenue from those taxpayers by audit. We especially consider how the nature of the tax authorities resource constraints shape the response to such adoption.

5.2.1 Empirical Motivation

We begin with a simple empirical fact to motivate our simple model. This fact comes from taxpayers' contesting auditors' assessments, which we observe in the operational audit data. The tax assessed by an auditor is their professional determination of the tax due given the taxpayers circumstances and the applicable tax laws, regulations, and revenue procedures. If the audited taxpayer (or their advisor) has a different interpretation of tax law, they can formally contest the assessment.[47] If the taxpayer contests the assessment, the IRS and the taxpayer enter into negotiations and, if these negotiations fail to reach

[47]The same is true of random audits. In both the NRP and our treatment of operational audits above, we use the initial assessment of the auditor, before any disagreement. The two measures considered in Figure 4 are therefore comparable in this respect.

an agreement, the case must be finally resolved in court. In complex circumstances where negotiations fail to come to an agreement quickly, the resulting litigation can take several years. Public data on such disagreements can be found in IRS (2020), Table 18.

Figure 10 depicts the share of assessed tax with which taxpayers disagree with their initial assessment (before negotiation), and the share of audited taxpayers who disagreed with their assessments. We observe that the share of tax dollars assessed that is subject to disagreement hovers around 25% through the bottom 90% of the income distribution, and then increase substantially in the top 10%, up to more than 60% in the very top bin. Individuals in the bottom bin, which includes those taxpayers with negative income, disagree at comparable rates to the very top bin, reflecting that audited individuals with negative reported income are typically high-wealth individuals. The share of taxpayers who disagree follows a similar pattern, but the overall share is significantly lower. That the dollar share is much larger than the taxpayer share implies that, perhaps unsurprisingly, the very largest assessments are typically the subject of disagreements.

We interpret this evidence as jointly informative about 1) the sophistication of avoidance/evasion strategies through the distribution, and 2) the extent to which taxpayers fight their assessments and attempt to negotiate. Legal experts and practitioners are well aware that audits of sophisticated high-wealth individuals can become quite complex and contentious. However the economic importance of these types of frictions for tax compliance and administration is not well-understood.[48] The magnitude of the differences in Figure 10 suggest that this is an important question. Perhaps the most salient implication of this fact for tax administration is that the high frequency of disputes and litigation makes recovering revenue from the top of the distribution via audit more costly. In the model below, we consider the implications of this notion for the allocation of audits through the distribution.

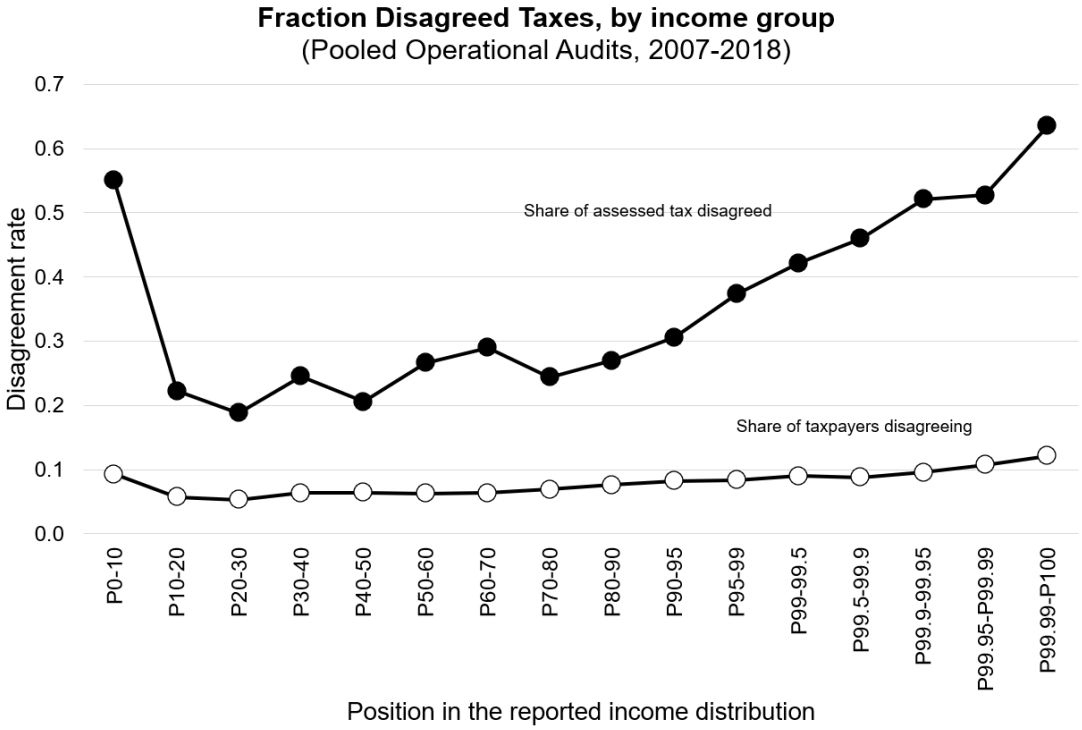
5.2.2 Model

Setup. There are two types of taxpayers, denoted by $\theta \in \{h, l\}$, which we basically think of as high- and low-income taxpayers. We first consider a revenue maximization problem with an exogenous resource constraint B. The tax authority decides how many of each type to audit which we denote by N_θ. Expected revenue raised by each type as a function of the number audited is $R_\theta(N_\theta)$. There is a constant marginal cost of auditing each type, c_θ. The objective is to maximize expected revenue net of costs.

The key difference between this model and the one we contrast it to later on is that we assume the total cost of audits cannot exceed some exogenous resource constraint B.

[48] See Blumenthal et al. (1998) for a model of audits as negotiations that may be relevant here.

FIGURE 10: CONTESTED AMOUNTS FOR OPERATIONAL AUDITS



Note: The top series of this figure plots the share of the total initial audit tax assessment that is contested by the taxpayer, across the income distribution. We rank taxpayers according to reported income in the tax year for which the taxpayer is under audit. The bottom series plots the share of audited taxpayers that contest their assessment amount. The data are pooled for fiscal years 2007-2018. The contested rate is very stable until the 90th percentile where it begins to increase and then rises sharply within the top 0.01% (up to 60%). The assessment share is significantly larger than the share of contesting taxpayers signifying that those with higher assessed values are more likely to contest. The large contested shares in the bottom of the distribution are mostly from taxpayers claiming large losses that are disallowed upon audit.

$$\max_{N_h, N_l} \quad R_h(N_h) + R_l(N_l) - c_h N_h - c_l N_l, \tag{3}$$

$$\text{subject to } c_h N_h + c_l N_l \leq B \tag{4}$$

Note that because of the presence of the resource constraint, this model is isomorphic to one in which the tax authority maximizes gross recovered revenue $R_h(N_h) + R_l(N_l)$ subject to the same resource constraint. The resource constraint requires that the last two terms in the objective function in equation (3) add up to a constant, so these terms become irrelevant for optimization.

This problem differs from an "optimal tax systems" approach to this question (Slemrod and Yitzhaki, 2002; Keen and Slemrod, 2017), in two important ways. Most importantly, the tax authority is given an exogenous resource constraint rather than simply maximizing net revenue, which we relax later. Addition-

ally, for simplicity, we do not account for distortions induced by changes in audit policy that can cause the optimal policy to deviate from revenue maximization, such as compliance costs. Accounting for such distortions would not change the main result of interest here.

The first-order condition for an interior optimum of this problem is

$$\frac{R_h'(N_h)}{c_h} = \frac{R_l'(N_l)}{c_l}. \tag{5}$$

Proposition 3. ***Comparative Statics of the Resource-Constrained Model.*** *In the optimization problem described by equation* (3)*,*

- $\frac{\partial N_h}{\partial c_h} < 0$
- $\frac{\partial N_l}{\partial c_h} > 0$ if and only if $-N_h R_h'' / R_h' > 1$.

That increasing c_h decreases N_h is unsurprising. More interesting is that in this model, the change in c_h has an effect on audits of low-income individuals. Because the tax authority is allocating finite resources to these two types of audits, the change in c_h has two effects on N_l, which are exactly analogous to an income and substitution effect in consumer choice theory. First, holding N_h fixed, increasing c_h leaves fewer resources available for audits of type l, which tends to decrease N_l: the income effect. Second, increasing c_h induces the tax authority to substitute toward auditing more type l taxpayers.

Which one of these effects dominates depends on the curvature of the revenue function for type h, $-N_h R_h'' / R_h'$, which determines whether total expenditure on h type audits goes up or down.

We next show that if we relax the exogeneity of the resource constraint, the spillover effect of an increase in c_h on audits of type l taxpayers disappears. Ignoring the resource constraints, the objective in (3) has simple first-order conditions that equate marginal revenue and marginal cost:

$$\begin{aligned} R_l' &= c_l \\ R_h' &= c_h. \end{aligned} \tag{6}$$

Proposition 4. ***Comparative Statics Without the Resource Constraint.*** *Consider the optimization problem described by equation* (3) *but ignore the resource constraint. In this model*

- $\frac{\partial N_h}{\partial c_h} < 0$
- $\frac{\partial N_l}{\partial c_h} = 0$.

Proposition 4 states that without an exogenous resource constraint, the spillover effects from an increase in c_h onto low-income types no longer occurs in this model.[49]

Contrasting Proposition 3 and 4 helps us understand how increased concealment effort by high-income taxpayers might affect low-income taxpayers, which we view as interesting given recent debates about the allocation of resources to various types of audits. The resource-constrained version of the model is closer to how tax administration works in the real world, where the IRS is given a budget by Congress and allocates these resources toward various types of enforcement. In this model, because the tax authority is devoting limited resources to all types of audits, increased concealment effort by high-income taxpayers can actually cause the tax authority to *substitute* toward auditing more low-income taxpayers, or it can deplete resources and cause fewer audits of low-income taxpayers. The unconstrained version of the model is closer in spirit to a model of optimal policy—subject to the caveats described e.g. by Slemrod and Yitzhaki (2002). The results for this version of the model model imply that increased concealment has no impact on *socially optimal* audit policy toward low-income individuals.

6 Conclusion

We find that substantial evasion at the top of the income distribution goes undetected in random audits. Investigating taxpayers who voluntarily declared hidden wealth or started reporting foreign bank accounts in 2009–2012 and who had been randomly audited just before, we find that in the vast majority of cases, the audits had failed to uncover offshore tax evasion. Focusing on taxpayers who earn business income through partnerships and S-corporations, we find that due to the resource constraints inherent to the conduct of random audits, a large fraction of this business income is not examined in the context of these audits, biasing detected evasion downward at the top.

Theoretically, we show that modelling the choice to conceal tax evasion from auditors can explain why random audits do not detect all evasion especially at the top. Empirically, we provide corrected estimates of the size and distribution of tax evasion in the United States. In our benchmark scenario, we find that under-reported income rises from about 7% of true income in the bottom 50% of the income distribution to 22% in the top 1%. Out of this 22%, about 6 percentage points correspond to sophisticated evasion that goes almost undetected in random audit. Accounting for tax evasion increases the top 1% income share in the United States.

It is important to note that random audit programs were not designed to estimate the tax gap for very

[49] Key for this result to obtain is that R_h does not depend on N_l and vice versa. This seems realistic, but it could be violated, for example, if auditing one type could lead to the discovery of information that is useful for auditing the other type.

high-income, high-wealth individuals. To experts who are familiar with these audits, our results may be unsurprising. However, we nevertheless view these results as important in light of an increased academic and policy interest in top income shares and tax evasion at the top. Countries around the world use random audits to estimate the tax gap (see, e.g., OECD, 2017, chapter 14). Our findings suggest that due to the limitations of random audits, more work is needed to estimate the extent of tax evasion at the top of the income distribution globally.

We stress that our estimates are likely to be conservative with regard to the overall amount of evasion at the top. From public reporting and anecdotal evidence, it seems likely that there are other specific forms of tax evasion that have the same properties as those we examine in this paper—sophistication and concentration among high income/wealth individuals. Such forms of evasion could include the abuse of syndicated conservation easements, micro-captive insurance schemes, private inurement in tax-exempt organizations, and the use of offshore trusts to evade tax. Many of these strategies involve pass-through businesses or other entities controlled by the taxpayer. The potential existence of many more such schemes underscores the main point of our theoretical results, that we should expect sophisticated evasion to be concentrated at the top of the income and wealth distribution. More research is needed to improve estimates of noncompliance at the very top in the United States.

We identify several other potentially fruitful avenues for future work. First, it would be valuable to consider the importance of sophisticated evasion and gray area avoidance strategies for optimal tax administration policies involving high-income, high-wealth taxpayers. Second, more research is needed to fully understand the gray area between avoidance and evasion, a line which can be blurry at the top of the income distribution and for large corporations. Future work could consider the implications of this notion for taxpayer behavior. Third, future research could consider strategic interaction between the tax authority and high-income individuals. We stopped short of such strategic, game-theoretic questions in our analysis, focusing separately on decisions by the individual taxpayer and by the tax authority. However, such strategic interactions may be empirically relevant and merit exploration in the future. Finally, future work could consider the implications of the ideas pursued here for white collar, financial crime more broadly, beyond sophisticated tax evasion.

References

Allingham, M. G. and Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of public economics*, 1(3-4):323–338.

Alstadaeter, A., Johannesen, N., and Zucman, G. (2018). Who owns the wealth in tax havens? Macro evidence and implications for global inequality. *Journal of Public Economics*, 162:89 – 100.

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

Andreoni, J., Erard, B., and Feinstein, J. (1998). Tax Compliance. *Journal of Economic Literature*, 36(2):818–860.

Auten, G. and Langetieg, P. (2020). The distribution of underreported income: What we can learn from the NRP? Presentation at the national tax association spring symposium on 14 may 2020.

Auten, G. and Splinter, D. (2019). Income inequality in the united states: Using tax data to measure long-term trends. Working paper.

Blumenthal, M., Christian, C., and Slemrod, J. (1998). The determinants of income tax compliance: Evidence from a controlled experiment in minnesota. NBER working paper no. 6575.

Brown, R. E. and Mazur, M. J. (2003). Irsâs comprehensive approach to compliance measurement. *National Tax Journal*, 56(3).

Carrillo, P., Pomeranz, D., and Singhal, M. (2017). Dodging the taxman: Firm misreporting and limits to tax enforcement. *American Economic Journal: Applied Economics*, 9(2):144–64.

Cooper, M., McClelland, J., Pearce, J., Prisinzano, R., Sullivan, J., Yagan, D., Zidar, O., and Zwick, E. (2016). Business in the United States: Who Owns it, and How Much Tax Do They Pay? *Tax Policy and the Economy*, 30(1):91–128.

Customs, H. R. . (2020). Measuring tax gaps 2020 edition: Tax gap estimates for 2018 to 2019. Hrmc offficial statistics release, 9 july 2020.

De Simone, L., Lester, R., and Markle, K. (2020). Transparency and tax evasion: Evidence from the foreign account tax compliance act (fatca). *Journal of Accounting Research*, 58(1):105–153.

DeBacker, J., Heim, B., Tran, A., and Yuskavage, A. (2020). Tax noncompliance and measures of income inequality. *Tax Notes*, February 17, 2020.

Feinstein, J. S. (1991). An Econometric Analysis of Income Tax Evasion and Its Detection. *The RAND Journal of Economics*, 22(1):14–35.

GAO (1995). Tax administration: Irs' partnership compliance activities could be improved. Letter report, 06/16/95, gao/ggd-95-151.

Guttentag, J. and Avi-Yonah, R. S. (2005). Closing the international tax gap.

IRS (2016). Federal Tax Compliance Research: Tax Gap Estimates for Tax Years 2008-2010. Publication 1415 (rev. 5-2016), Washington, DC.

IRS (2019). Federal Tax Compliance Research: Tax Gap Estimates for Tax Years 2011-2013. Publication 1415 (rev. 9-2019), Washington, DC.

IRS (2020). Internal revenue service data book 2019. Publication 55-b (rev. 06-2020), Washington, DC.

Johannesen, N., Langetieg, P., Reck, D., Risch, M., and Slemrod, J. (2020). Taxing hidden wealth: The consequences of us enforcement initiatives on evasive foreign accounts. *American Economic Journal: Economic Policy*, 12(3):312–46.

Johns, A. and Slemrod, J. (2010). The Distribution of Income Tax Noncompliance. *National Tax Journal*, 63(3):397–418.

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

Kleven, H. J., Knudsen, M. B., Kreiner, C. T., Pedersen, S., and Saez, E. (2011). Unwilling or Unable to Cheat? Evidence From a Tax Audit Experiment in Denmark. *Econometrica*, 79(3):pp. 651–692.

Kleven, H. J., Kreiner, C. T., and Saez, E. (2016). Why can modern governments tax so much? an agency model of firms as fiscal intermediaries. *Economica*, 83(330):219–246.

Luttmer, E. F. and Singhal, M. (2014). Tax morale. *Journal of economic perspectives*, 28(4):149–68.

OECD (2017). *Tax Administration 2017*.

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

Piketty, T., Saez, E., and Zucman, G. (2018). Distributional national accounts: methods and estimates for the united states. *The Quarterly Journal of Economics*, 133(2):553–609.

Saez, E. and Zucman, G. (2016). Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data. *The Quarterly Journal of Economics*, 131(2):519–578.

Saez, E. and Zucman, G. (2020). Trends in US Income and Wealth Inequality: Revising after the Revisionists. NBER working paper no. 27921.

Sarin, N. and Summers, L. H. (2020). Understanding the revenue potential of tax compliance investment. NBER working paper no. 27571.

Slemrod, J., Collins, B., Hoopes, J. L., Reck, D., and Sebastiani, M. (2017). Does credit-card information reporting improve small-business tax compliance? *Journal of Public Economics*, 149:1–19.

Slemrod, J. and Yitzhaki, S. (2002). Tax avoidance, evasion, and administration. In *Handbook of public economics*, volume 3, pages 1423–1470. Elsevier.

Smith, M., Zidar, O., and Zwick, E. (2019). Top wealth in the united states: New estimates and implications for taxing the rich. *Unpublished*.

TIGTA (2015). Improvements are needed in resource allocation and management controls for audits of high-income taxpayers. Treasury inspector general for tax administration publication 2015-30-078.

TIGTA (2020). High-income nonfilers owing billions of dollars are not being worked by the internal revenue service. Treasury inspector general for tax administration publication 2020-30-015.

United States Senate (2008). Tax Haven Banks and U.S. Tax compliance. Staff report of the permanent subcommittee on investigations, Washington, DC.

United States Senate (2014). Offshore Tax Evasion: The Effort to Collect Unpaid Taxes on Billions in Hidden Offshore Accounts. Technical report, Washington, DC.

Yitzhaki, S. (1974). A note on 'income tax evasion: A theoretical analysis'. *Journal of Public Economics*, 3(2):201–202.

Zucman, G. (2013). The Missing Wealth of Nations: Are Europe and the U.S. net Debtors or net Creditors? *The Quarterly Journal of Economics*, 128(3):1321–1364.

Zucman, G. (2014). Taxing Across Borders: Tracking Personal Wealth and Corporate Profits. *Journal of Economic Perspectives*, 28(4):121–148.

A Additional Results

TABLE A1: TAX EVASION DETECTED IN NRP RANDOM AUDITS WITHOUT DCE CORRECTION: DECOMPOSITION BY INCOME TYPE

	Full Population				Top 1%			
	Total income of this type/ Total income (%)	Total under-reported income of this type/ Total under-reported income (%)	Total under-reported income of this type/ Total income (%)	Total under-reported income of this type/ Total income of this type (%)	Total income of this type/ Total income (%)	Total under-reported income of this type/ Total under-reported income (%)	Total under-reported income of this type/ Total income (%)	Total under-reported income of this type/ Total income of this type (%)
Capital Gains	5.6	7.0	0.27	4.9	21.0	16.7	0.43	2.0
Dividends	3.7	2.7	0.10	2.8	8.5	3.6	0.09	1.1
Interest	1.8	0.7	0.03	1.4	3.0	1.8	0.05	1.6
Line 21 Other Income	1.0	10.3	0.40	41.1	2.5	7.5	0.19	7.5
Partnerships and S Corp	5.9	6.4	0.25	4.2	21.5	16.7	0.43	2.0
Rental	0.7	8.6	0.33	46.4	1.6	4.9	0.13	8.1
Schedule C	5.3	50.2	1.95	36.7	4.5	41.9	1.07	23.5
Wages	70.8	3.7	0.14	0.2	38.4	2.7	0.07	0.2
Other	5.2	10.4	0.40	0.1	-1.0	4.1	0.11	-0.1
Total	100.0	100.0	3.87		100.0	100.0	2.45	

Note: This table describes the composition of detected under-reported income in the 2006–2013 NRP data, before any correction for undetected noncompliance (in particular before any DCE correction). The NRP shares of each type of income in total income are similar to income shares we observe in SOI data, but the NRP shares are built using corrected income here. Consequently the largest differences with SOI income shares are observed for types of income with significant detected evasion. Note that "Form 1040 Other Income" in Figure 2 is referred to as "Line 21 Other Income" here, as this item appears on Line 21 of the Form 1040, while the residual "Other" category in the penultimate row refers to all other components of income.

TABLE A2: TAX EVASION DETECTED IN NRP RANDOM AUDITS WITH DCE CORRECTIONS: DECOMPOSITION BY INCOME TYPE

	Full Population				Top 1%			
	Total income of this type/ Total income (%)	Total under-reported income of this type/ Total under-reported income (%)	Total under-reported income of this type/ Total income (%)	Total under-reported income of this type/ Total income of this type (%)	Total income of this type/ Total income (%)	Total under-reported income of this type/ Total under-reported income (%)	Total under-reported income of this type/ Total income (%)	Total under-reported income of this type/ Total income of this type (%)
Capital Gains	5.5	5.2	0.73	13.1	18.5	7.4	1.44	7.8
Dividends	3.5	1.8	0.24	6.9	7.3	1.7	0.33	4.5
Interest	1.7	0.7	0.09	5.5	2.6	0.6	0.12	4.6
Line 21 Income	2.3	13.6	1.89	83.1	5.2	20.5	3.95	75.8
Partnerships and S Corp	6.3	9.1	1.26	20.0	20.3	16.1	3.11	15.3
Rental	1.7	9.8	1.37	80.0	2.2	5.5	1.06	47.1
Schedule C	9.4	46.1	6.40	68.1	10.6	38.2	7.38	69.7
Wages	64.0	4.1	0.56	0.9	32.9	2.2	0.42	1.3
Other	5.5	9.7	1.35	0.2	0.4	7.8	1.50	3.7
Total	100.0	100.0	13.89		100.0	100.0	17.81	

Note: This table describes the composition of detected under-reported income in the 2006–2013 NRP data, before any correction for undetected noncompliance (in particular before any DCE correction). The NRP shares of each type of income in total income are similar to income shares we observe in SOI data, but the NRP shares are built using DCE-adjusted income here. Consequently the largest differences with SOI income shares are observed for types of income with significant detected evasion. Note that "Form 1040 Other Income" in Figure 2 is referred to as "Line 21 Other Income" here, as this item appears on Line 21 of the Form 1040, while the residual "Other" category in the penultimate row refers to all other components of income.

TABLE A3: ESTIMATED TAX DUE AND INCOME EVADED, IN BILLION DOLLARS, NRP AND OFFSHORE

	NRP income under-reported	NRP tax evaded	Offshore income evaded	Offshore tax due
Bottom 99%	977.9	82.9	6.6	1.0
Top 1% (inclusive)	144.0	14.3	53.7	11.6
Top 0.1% (inclusive)	42.1	4.0	39.9	8.7
Top 0.01%	7.7	0.3	22.5	5.3
All	1,121.9	97.2	60.3	12.6

Note: This table reports estimates of the total income under-reported and taxes evaded, for NRP results with no DCE correction and for offshore estimates. All the figures are presented in billion dollars.

TABLE A4: INCOME SHARES IN 2006-2013 NRP DATA AND IN 2001 NRP DATA

	2006–2013 Before exam	2006–2013 After exam No DCE	2006–2013 After exam With DCE	2006–2013 Our benchmark	2001 Before exam	2001 After exam With DCE
P0-10	-2.3	-1.9	-1.5	-1.4	0.1	0.3
P10-20	1.3	1.2	1.0	1.0	1.6	1.6
P20-30	2.2	2.0	1.7	1.7	2.7	2.7
P30-40	3.0	2.9	2.5	2.4	3.9	3.9
P40-50	4.3	4.2	3.6	3.6	5.2	5.2
P50-60	5.9	5.8	5.0	5.0	6.8	6.7
P60-70	8.4	8.2	7.3	7.2	8.9	8.8
P70-80	11.8	11.7	10.8	10.7	11.7	11.5
P80-90	16.8	16.9	16.2	16.0	16.0	15.6
P90-95	11.9	12.1	12.2	12.1	11.0	10.9
P95-99	16.5	16.8	19.0	18.9	14.4	14.9
P99-99.5	4.2	4.3	5.2	5.2	3.7	3.8
Top 0.5%	16.1	15.8	17.0	17.6	14.1	14.0

Note: This table reports the distribution of income in the 2006–2013 NRP data studied in this paper and in the 2001 NRP data as reported in Johns and Slemrod (2010, Table 5). Tax units are ranked by their estimated true income (equal to reported income plus estimated under-reported income). Income is Adjusted Gross Income (AGI) in Johns and Slemrod (2010) and market income in our series (defined as total income reported on form 1040 minus Social Security benefits, unemployment insurance benefits, alimony, and state refunds). Series in columns 3, 4, and 6 all use the same DCE methodology.

TABLE A5: INCOME AND TAXES OWED: BEFORE VS. AFTER ACCOUNTING FOR TAX EVASION (BILLIONS OF $2012)

	Income				Tax			
	Reported	Corrected	DCE-corrected	Our benchmark	Reported	Corrected	DCE-corrected	Our benchmark
P0-P90	4,181	4,314	4,404	4,407	362	407	415	415
P90-P95	965	1,024	1,153	1,153	126	140	156	157
P95-P99	1,342	1,423	1,793	1,799	238	258	324	327
P99-P99.5	339	360	490	492	79	86	114	116
P99.5-P99.9	562	582	733	756	140	146	187	196
P99.9-P100	747	762	876	925	180	185	217	233
Total	8,136	8,464	9,449	9,531	1,126	1,222	1,414	1,444

Note: This table reports aggregate income and taxes by income group before vs. after correction for tax evasion, as estimated in the NRP (cols. 1–3 and 5–7) and in our benchmark that factors in sophisticated evasion (cols. 4 and 8). Numbers are in billions of 2012 dollars and correspond to annual averages over the period 2006-2013. The table shows that the standard federal income tax gap (i.e., DCE-corrected taxes owed minus taxes paid) is $1,414 - $1,126 = $288 billion per year over 2006-2013. In our benchmark estimates, the tax gap is $1,444 - $1,126 = $318 billion per year over that period, with almost all the difference with the standard estimate coming from the top 0.5%. Our correction for sophisticated evasion increases the aggregate tax gap by a factor of 1.1 on aggregate, but by a factor of 1.4 for the top 0.1% (and 1.7 for the top 0.01%, not separately reported in this table). This should be seen as a conservative correction, given that it only factors in two forms of sophisticated evasion (offshore and pass-through businesses). We note that the change in the tax gap ($30 billion) and the change in income under-reporting gap ($82 billion) imply a marginal tax rate on our adjustments of 37%, higher that the top marginal income tax rate of 35% in 2006–2012. This occurs because when we remove 57% of all DCE-adjusted pass-through income under-reporting to avoid double counting and then add our benchmark total pass-through evasion, we effectively relocate some under-reported income from the bottom of the distribution to the top, which slightly increases the tax due on income under-reporting that was already accounted for in the DCE-adjusted total.

TABLE A6: CHANGE IN INCOME AND TAXES OWED: BEFORE VS. AFTER ACCOUNTING FOR TAX EVASION (BILLIONS OF $2012)

	Income				Tax			
	After exam − reported	DCE-corrected − after exam	Sophisticated − after exam	Our benchmark − DCE-corrected	After exam − reported	DCE-corrected − after exam	Sophisticated − after exam	Our benchmark − DCE-corrected
P0-P90	133	90	7	3	45	8	1	0
P90-P95	59	128	6	0	14	16	2	0
P95-P99	80	371	28	5	20	66	9	4
P99-P99.5	21	130	16	1	7	28	5	2
P99.5-P99.9	20	151	41	23	6	42	13	8
P99.9-P100	14	114	63	49	5	32	19	16
Total	328	985	161	82	96	192	50	30

Note: This table reports the change in aggregate income and taxes by income group (i) when correcting reported incomes and taxes in the NRP without DCE adjustment (cols. 1 and 5), (ii) when adding the DCE adjustment to exam-corrected NRP data (cols. 2 an 6) (ii) when adding sophisticated evasion to exam-corrected NRP data before DCE adjustment (cols. 3 and 7), and (iii) in our benchmark scenario that adds sophisticated evasion to the DCE-adjusted NRP after having removed 57% of DCE-adjusted pass-through business income evasion (cols. 4 and 8). Numbers are in billions of 2012 dollars and correspond to annual averages over the period 2006-2013. See also notes to Table A5.

TABLE A7: SHARES OF TAXES PAID AND UNPAID, 2006-2013, IN % OF TOTAL TAXES PAID OR UNPAID

	Taxes paid	Taxes unpaid NRP with DCE	Taxes unpaid Our benchmark
P0-10	0.2	0.2	0.2
P10-20	0.4	0.2	0.2
P20-30	0.7	0.6	0.5
P30-40	1.0	0.9	0.8
P40-50	1.7	1.5	1.3
P50-60	2.9	2.3	2.0
P60-70	4.8	3.6	3.2
P70-80	7.9	6.2	5.6
P80-90	12.7	10.3	9.3
P90-95	11.2	10.4	9.6
P95-99	21.2	26.8	25.4
P99-99.5	7.0	10.7	10.3
P99.5-P99.9	12.4	14.9	16.2
P99.9-99.95	3.7	4.1	4.8
P99.95-P99.99	5.9	4.3	5.6
P99.99-100	6.5	3.0	5.0
Top 1%	35.5	37.1	41.8

Note: This table reports the distribution of federal individual income taxes paid and unpaid taxes (for different measures of unpaid taxes) on average over 2006–2013. Tax units are ranked by their reported income in the first column and their estimated true income in the second and third column. The first column shows the distribution of paid taxes. The second column shows the distribution of unpaid taxes using the DCE-adjusted NRP random audits. The last column shows the distribution of unpaid taxes in our benchmark scenario (described in Section 4.1).

FIGURE A1: TAX EVADED (% OF TAX OWED) IN RANDOM AUDITS WITHOUT DCE CORRECTION

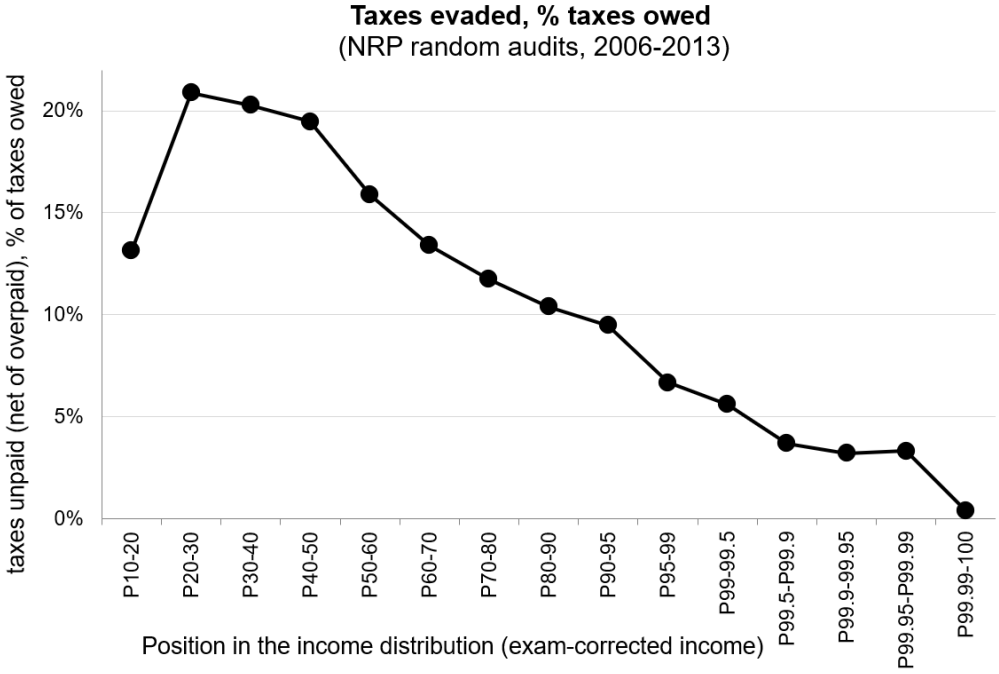


FIGURE A2: TAX EVADED (% OF TAX OWED) IN RANDOM AUDITS WITHOUT DCE CORRECTION, EXCLUDING TAX CREDITS

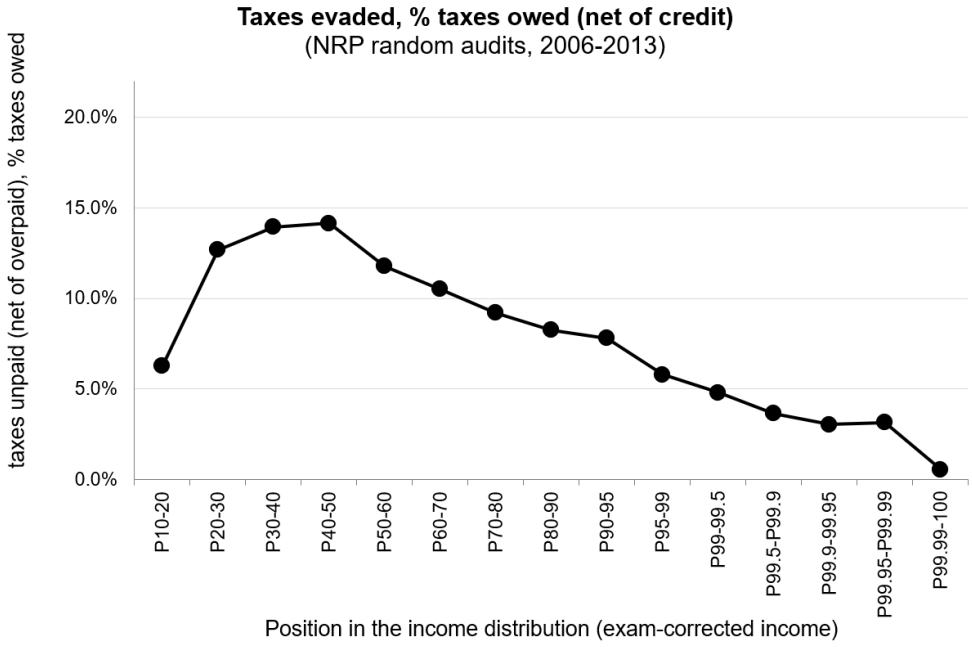
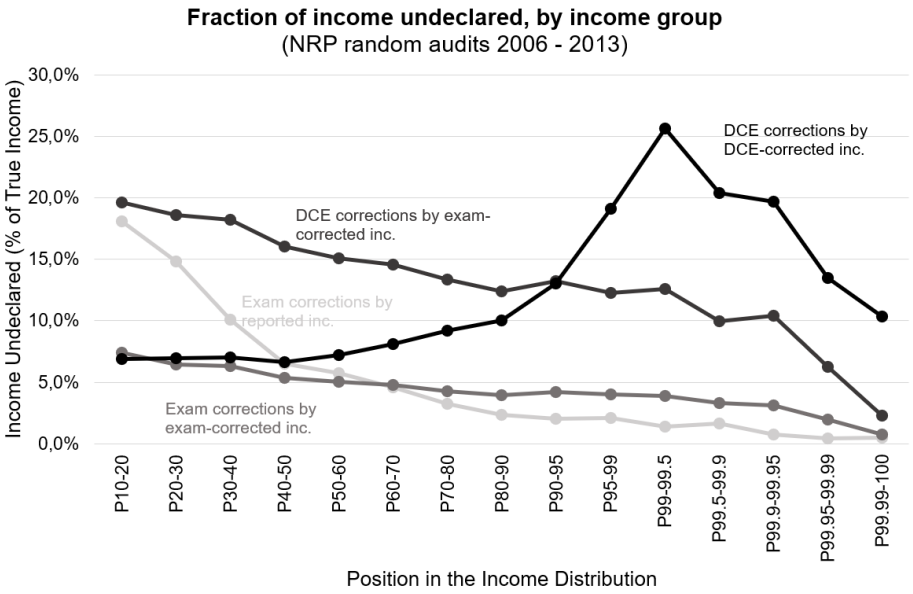
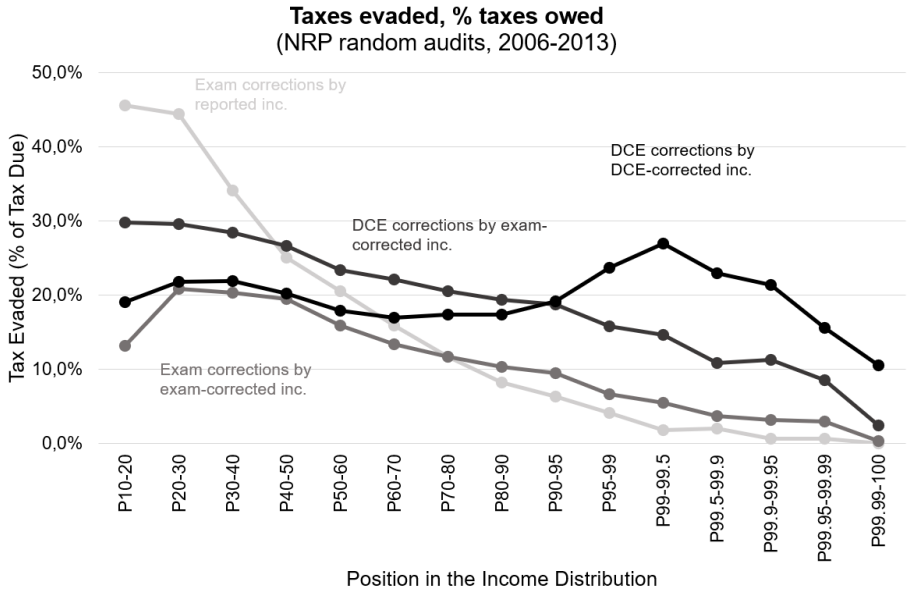


FIGURE A3: THE INFLUENCE OF RE-RANKING AND DCE ON ESTIMATED RATES OF EVASION

(a) Income Under-Reporting



(b) Tax Evaded



Note: This figure illustrates the impact of reranking and DCE adjustment on the profile of income under-reporting and the tax gap through the income distribution. We begin with "Exam corrections by reported income," which ranks taxpayers by originally reported income and calculates income and under-reporting gaps using exam corrections only (i.e. no DCE). We then continue to use exam corrections only but re-rank individuals by exam-corrected income in "exam corrections by exam-corrected income," which matches Figure 2. We find that this re-ranking substantially decreases estimated rates of evasion in the bottom 50% of the distribution. Third we implement DCE corrections to estimate rates of evasion, but for illustrative purposes we continue to rank individuals by their exam-corrected income, in "DCE corrections by exam-corrected income." Finally, we re-rank by DCE corrected income in "DCE corrections by DCE-corrected income." From comparing the third to fourth step we find that re-ranking substantially increases the rate of evasion in the top 5% of the distribution, and this re-ranking drives the increasing profile of evasion from DCE corrections in Figure 3.

FIGURE A4: TAX EVADED (% OF TAX OWED) IN RANDOM AUDITS WITH VS. WITHOUT DCE CORRECTION

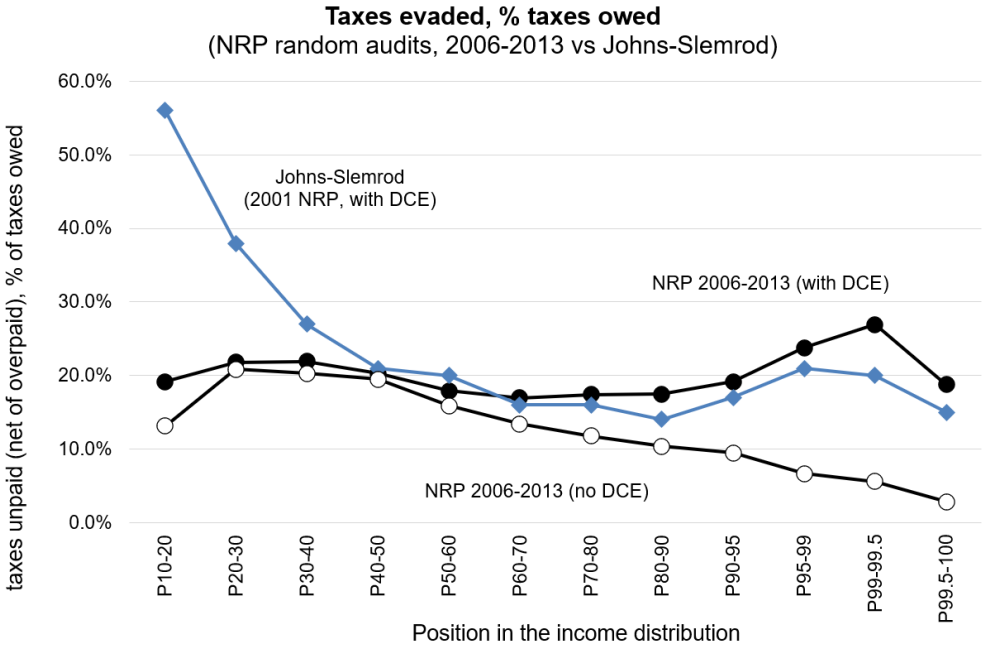
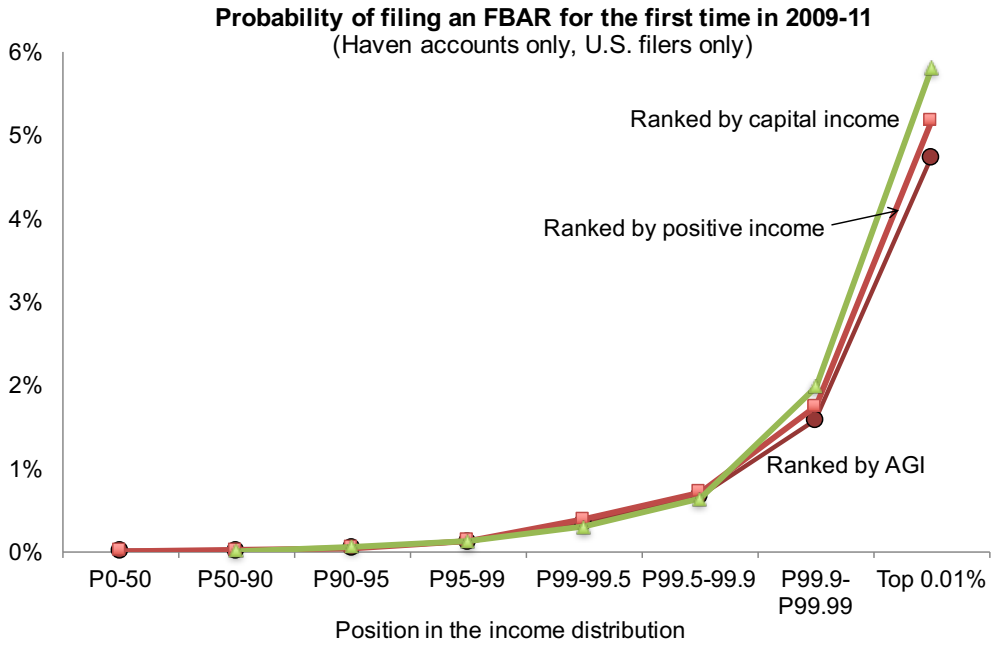


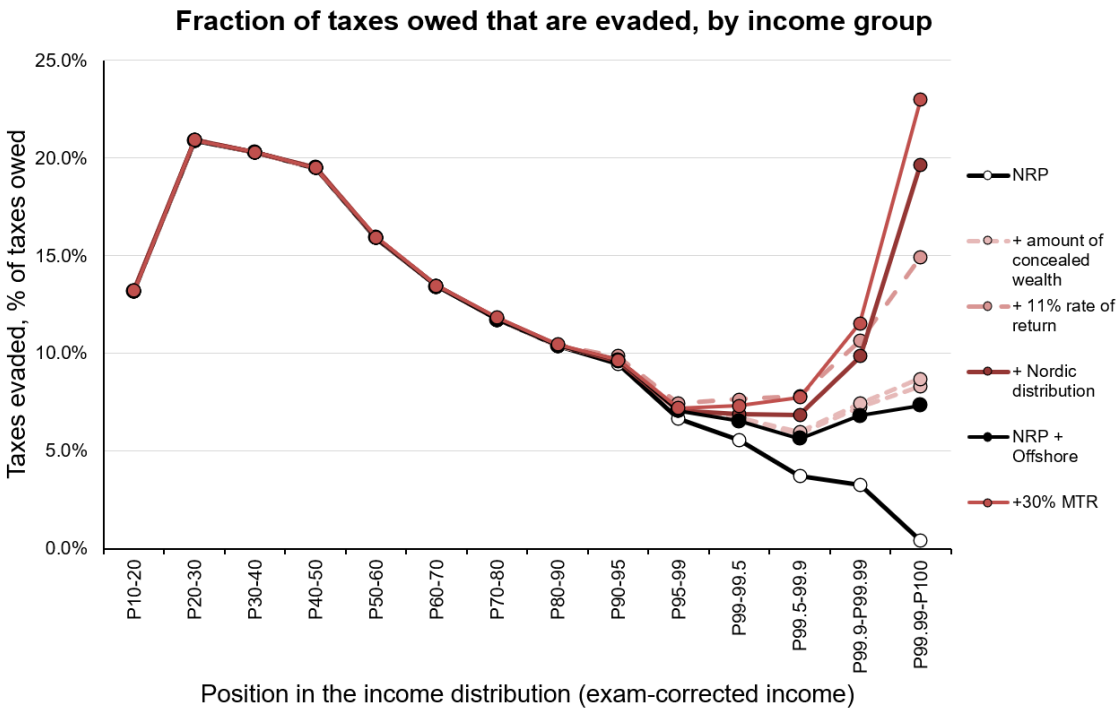
FIGURE A5



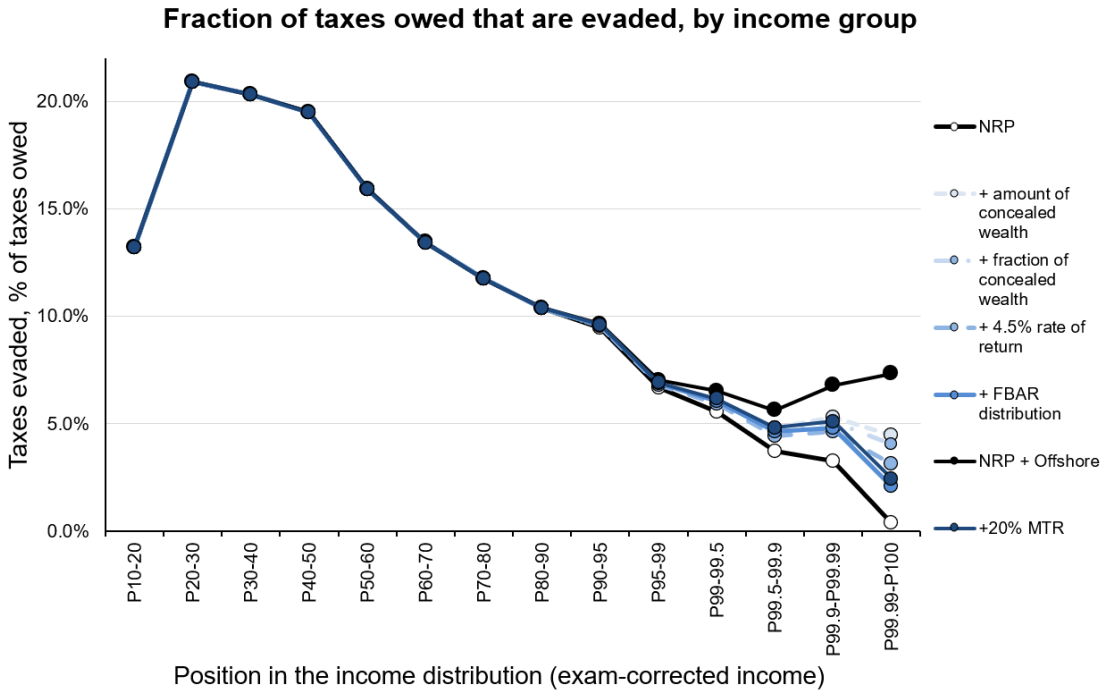
Note: This figure plots the fraction of the population within each part of the income distribution that are present in the first-time FBAR filer sample. We observe that the probability of being in the sample is much higher at the very top of the income distribution, with a nearly trivial fraction of the bottom 99 percent of the income distribution disclosing an offshore account. We observe that the overall profile is very similar for the three different income concepts, though it is steepest for capital income, followed by positive income.

FIGURE A6: DECOMPOSING SENSITIVITY ANALYSIS FOR OFFSHORE WEALTH

(a) Upper Bound

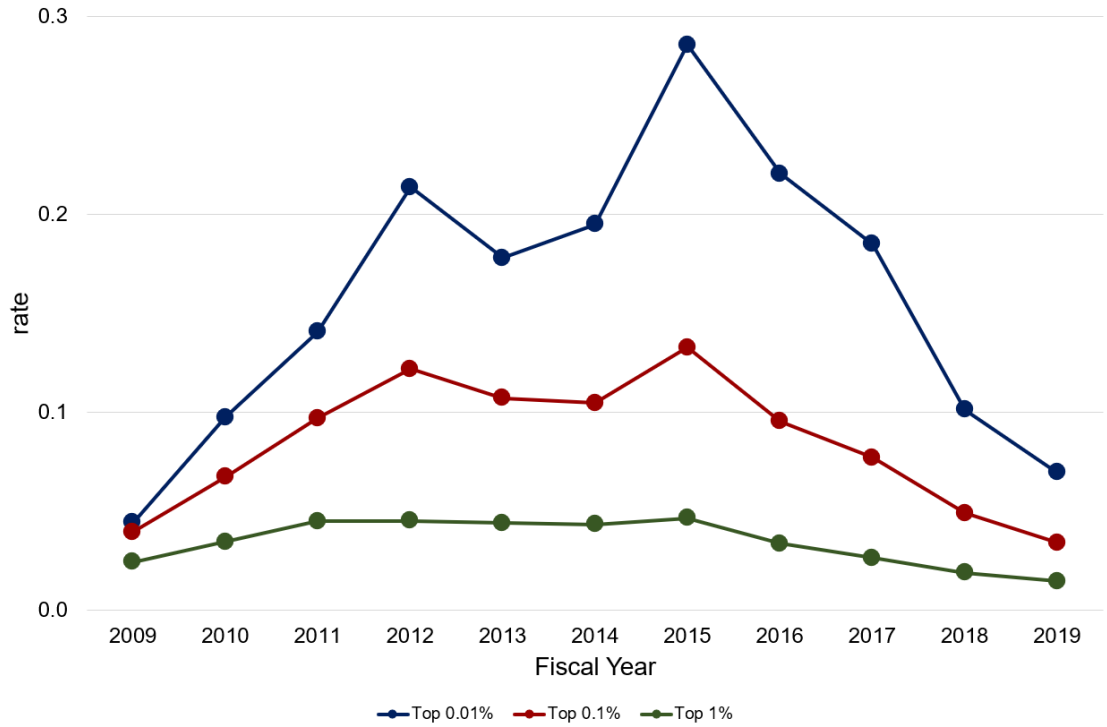


(b) Lower Bound



Note: This figure plots taxes evaded over taxes due by rank in the income distribution with and without accounting for offshore wealth. Taxpayers are ranked by their estimated true AGI. In either figure, we begin with our preferred scenario for offshore wealth and then progressively add assumptions for the alternative scenarios described in Table 1.

FIGURE A7: FRACTION OF INDIVIDUALS AUDITED BY YEAR



Note: This figure plots audit rates over time in the three income groups. We observe that audit rates are highest at the very top, and they increase and then decline through the period of observation period.

FIGURE A8: TOTAL ASSESSED TAXES : OPERATIONAL AUDITS VERSUS NRP RANDOM AUDITS

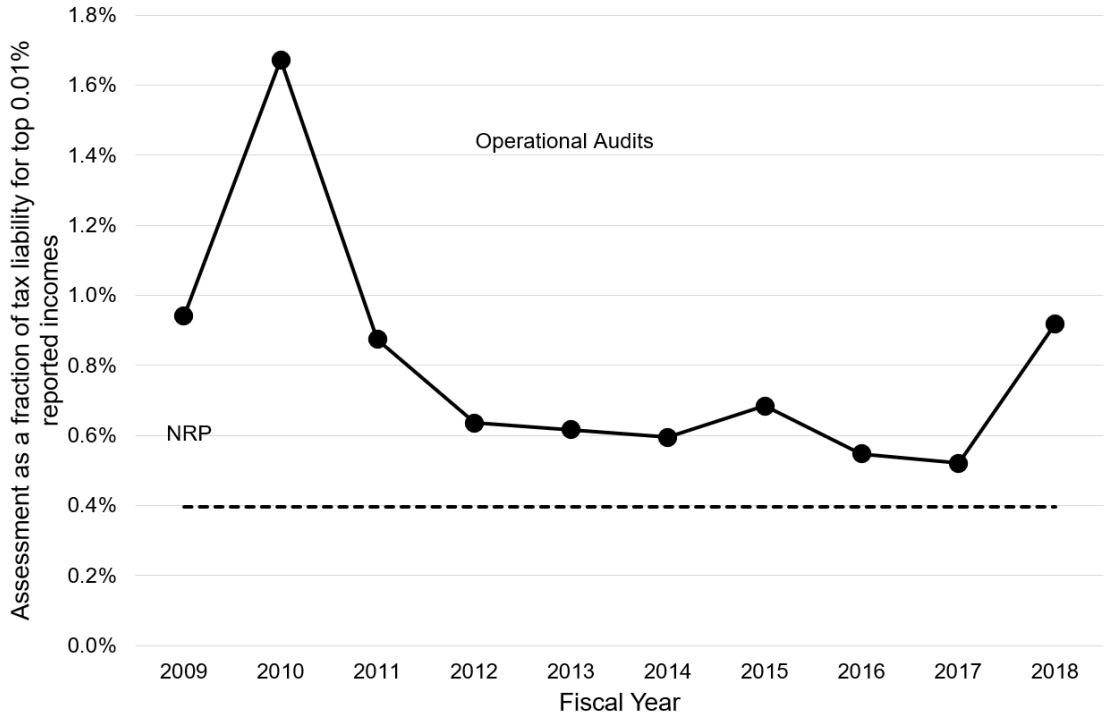


FIGURE A9

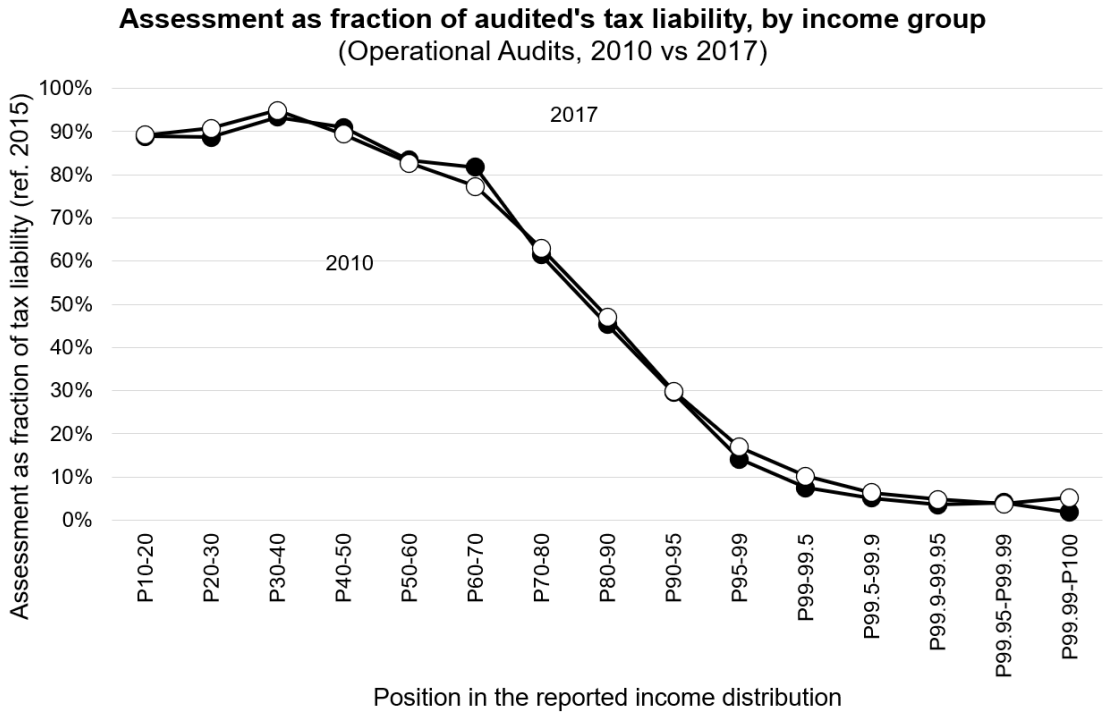


FIGURE A10: TAX GAP: THE EFFECT OF ACCOUNTING FOR UNREPORTED OFFSHORE INCOME

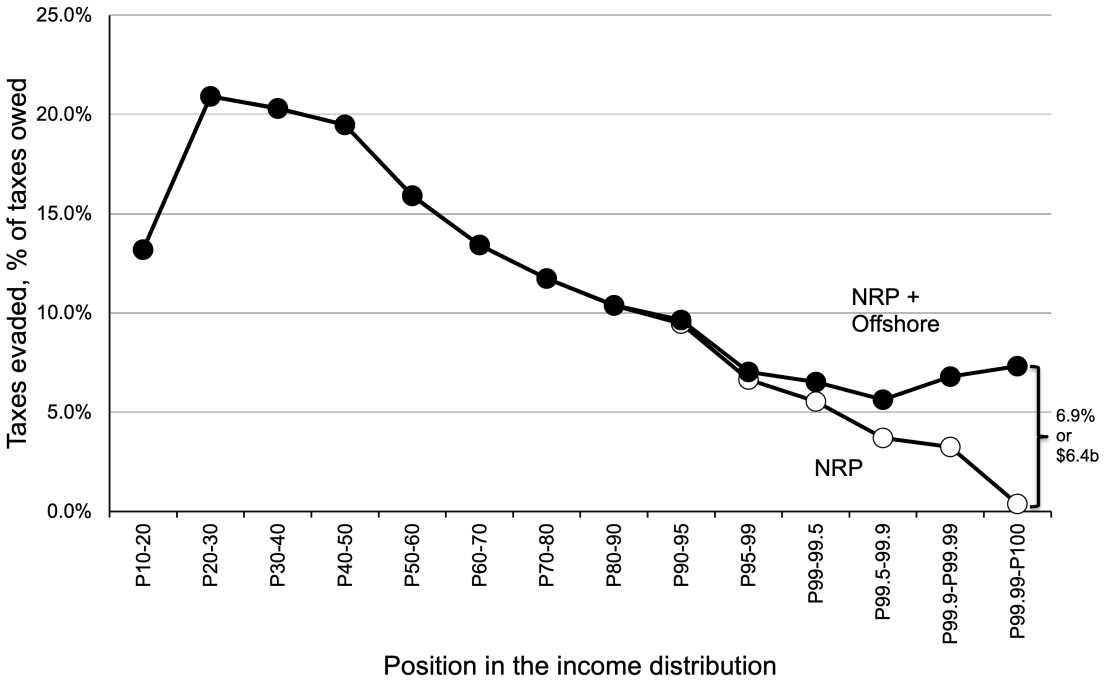


FIGURE A11: TAX GAP: THE EFFECT OF ACCOUNTING FOR UNREPORTED OFFSHORE INCOME (SENSITIVITY ANALYSIS)

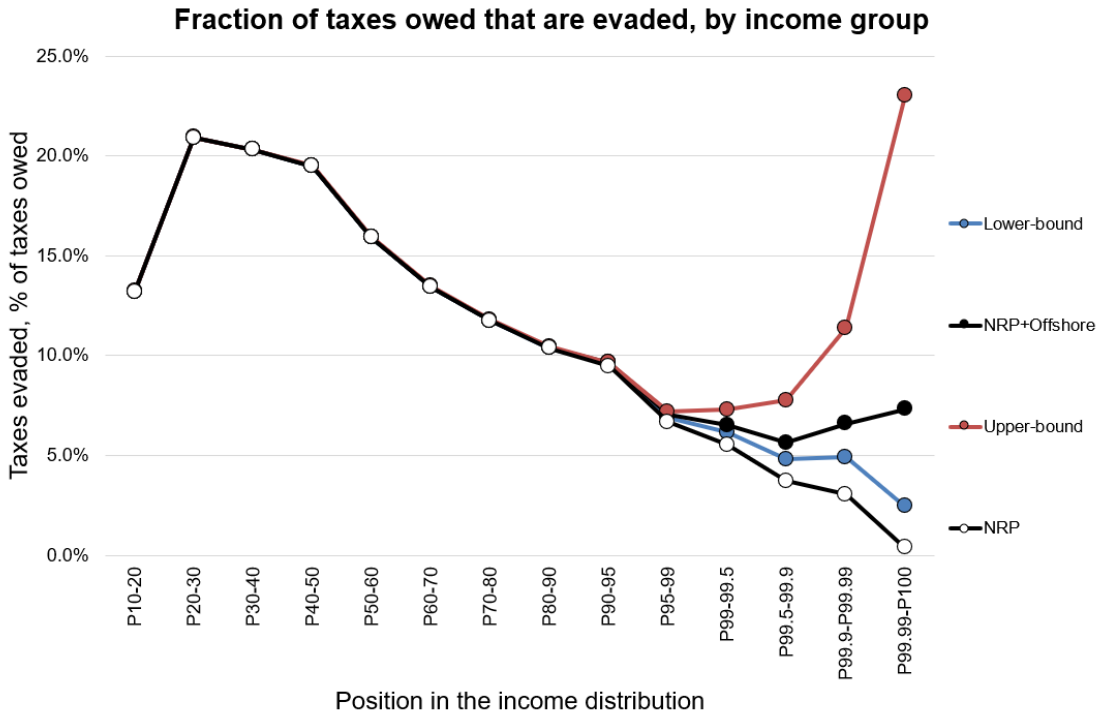


FIGURE A12: UNDER-REPORTED INCOME: THE EFFECT OF ACCOUNTING FOR PASS-THROUGH BUSINESS LOSSES

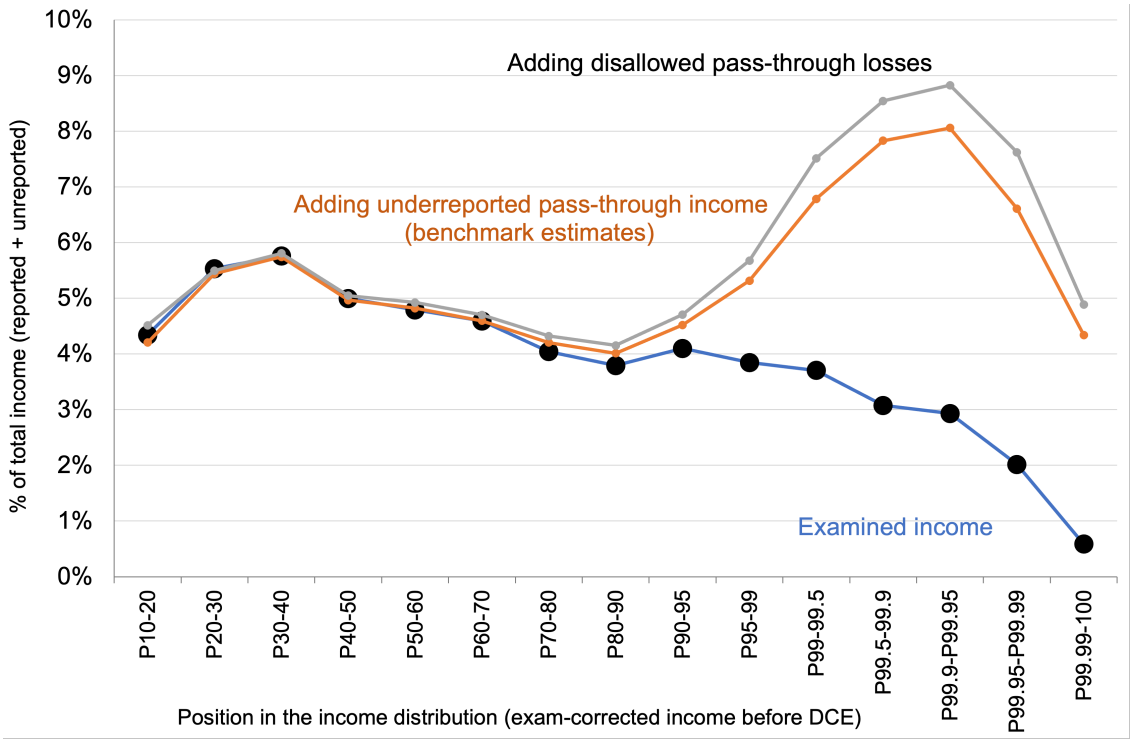


FIGURE A13: UNDER-REPORTED INCOME: THE EFFECT OF ACCOUNTING FOR CIRCULAR PARTNERSHIPS

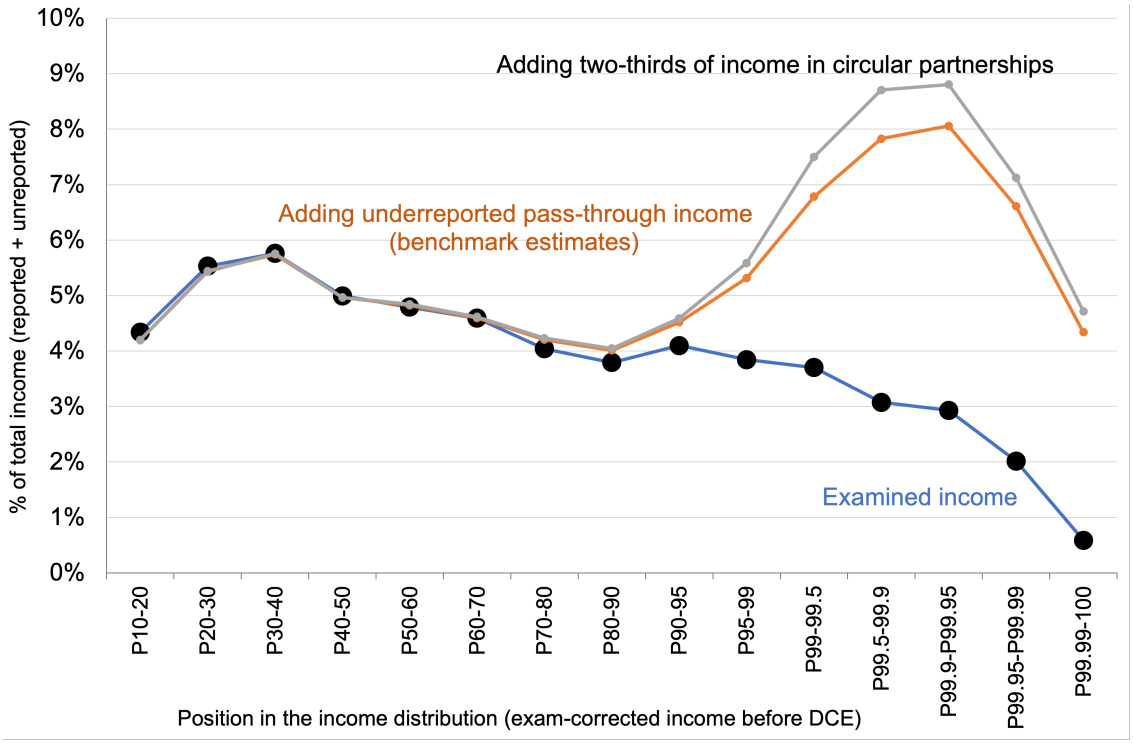


FIGURE A14: UNDER-REPORTED INCOME: THE EFFECT OF ACCOUNTING FOR PASS-THROUGH BUSINESS INCOME, SENSITIVITY ANALYSIS

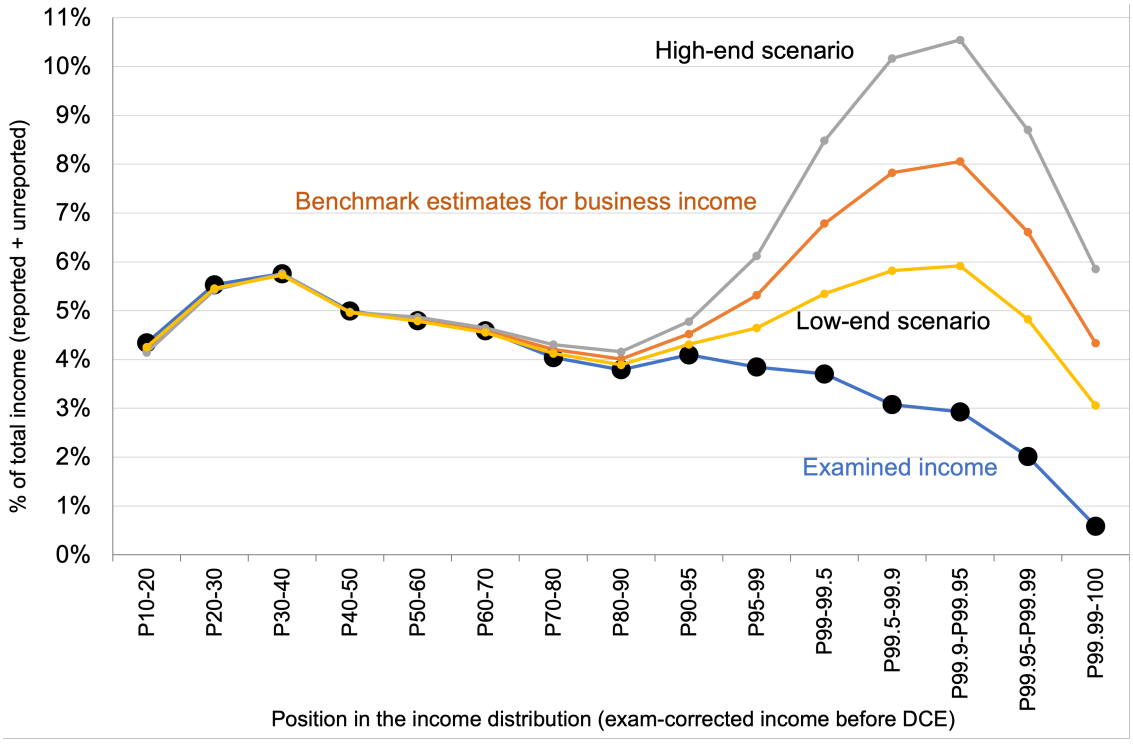


FIGURE A15: UNDER-REPORTED INCOME: THE EFFECT OF ACCOUNTING FOR PASS-THROUGH INVESTMENT INCOME, SENSITIVITY ANALYSIS

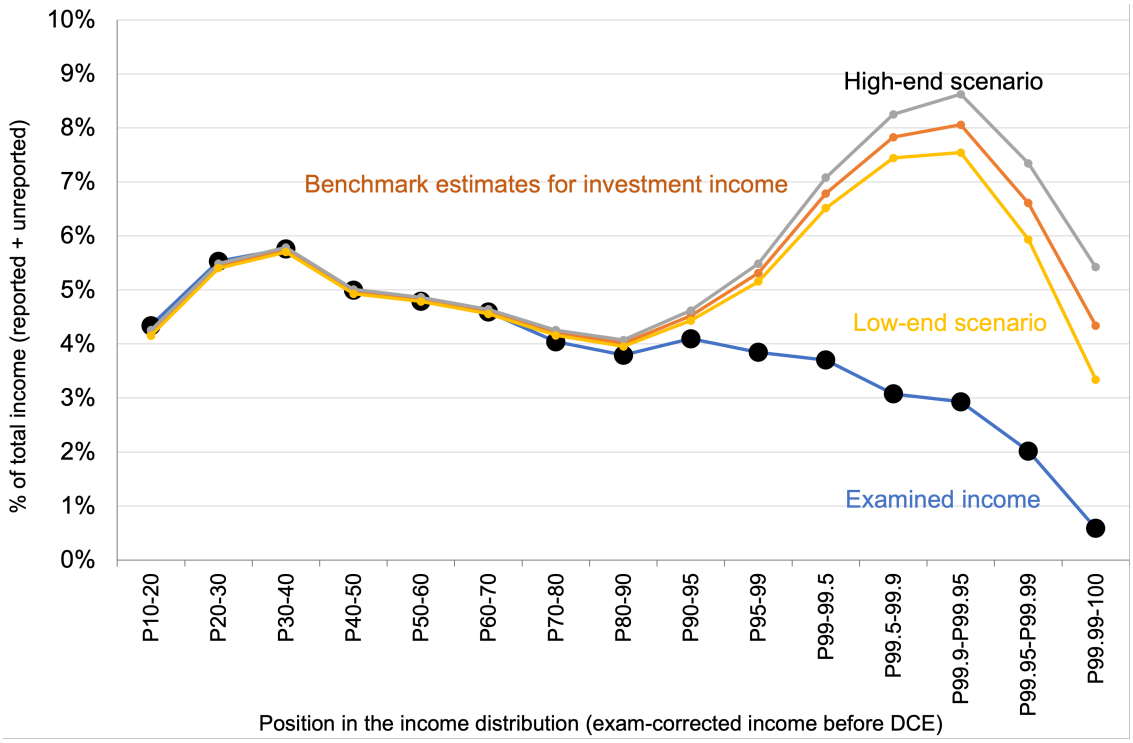


FIGURE A16: UNPAID TAXES (% OF TAXES OWED): BENCHMARK ESTIMATES

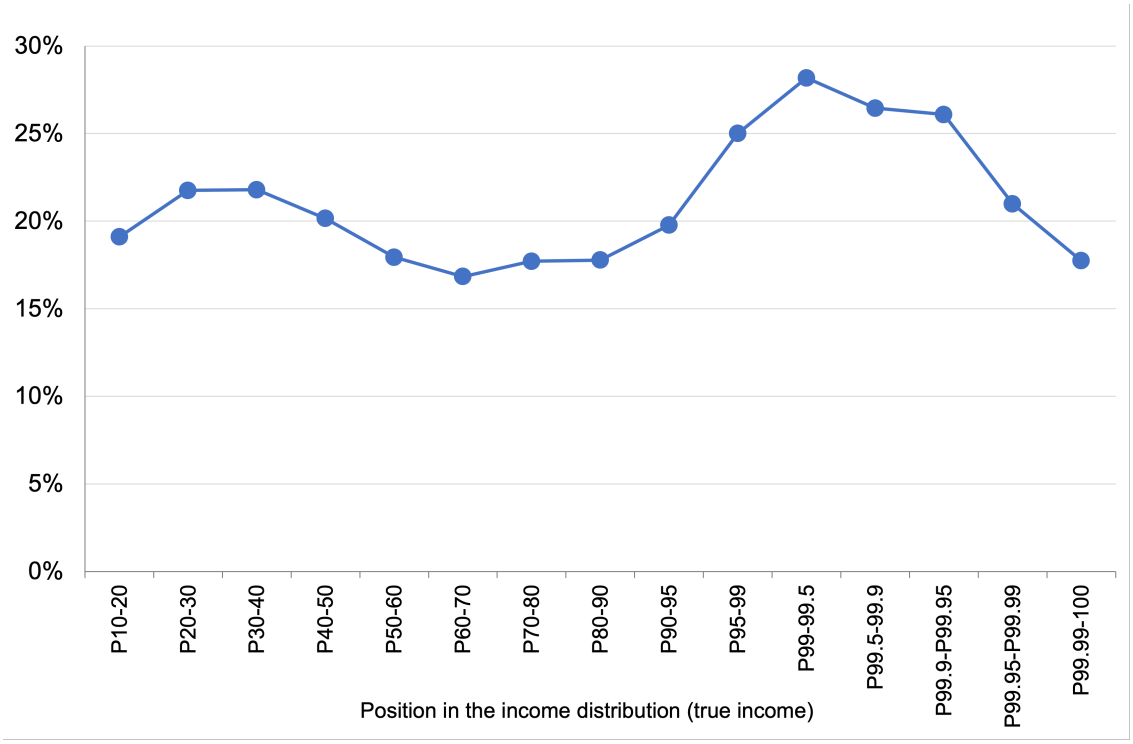
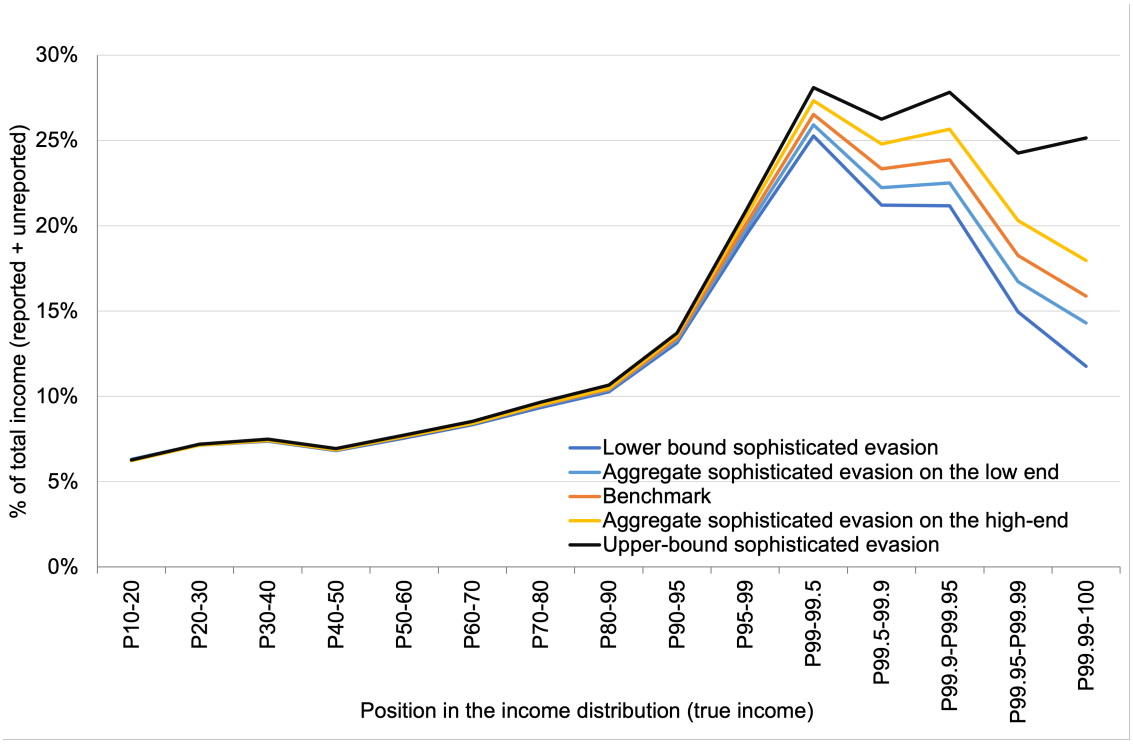


FIGURE A17: UNREPORTED INCOME (% OF TRUE INCOME): SENSITIVITY ANALYSIS



Note: This figure plots the results of the different scenarios described in Section 4.1

B Proofs

Lemma 2. *Under Assumption 1, as y becomes arbitrarily large, $g_1(y, p_1) - g_0(y, p_1)$ converges to zero.*

Proof. We can re-express the optimization problem with (g, a) as the choice variables rather than (e, a). We denote the fixed cost κ as a share of income by $\tilde{\kappa} = k/y$. We can then express consumption in the detected and undetected state by $c_D = (1 - \tau - \theta\tau g - \tilde{\kappa}a)y$, and $c_N = (1 - \tau + \tau g - \tilde{\kappa}a)y$. respectively. The first order condition with respect to g of the optimization problem in equation 1 is

$$\frac{u'(c_D)}{u'(c_N)} = \frac{1-p}{p\theta}. \tag{7}$$

We wish to compare $g_1(y, p_1)$ and $g_0(y, p_1)$ at large y. As both of these are evaluated at $p = p_1$, the right-hand side of (7) is constant for this comparison. Comparing the first order conditions under $a = 1$ and $a = 0$, we have:

$$\frac{u'((1 - \tau - \theta\tau g_1(y, p_1) - \tilde{\kappa})y)}{u'((1 - \tau + \tau g_1(y, p_1) - \tilde{\kappa})y)} = \frac{u'((1 - \tau - \theta\tau g_0(y, p_1))y)}{u'((1 - \tau + \tau g_0(y, p_1))y)}. \tag{8}$$

As y becomes large, $\tilde{\kappa} = \kappa/y$ becomes arbitrarily small. As $u'' < 0$, the LHS and RHS of equation (8) are invertible in e. Finally, by Assumption 1, for arbitrarily large y, $g_0(y, p_1)$ on the RHS converges to a strictly positive constant. Altogether, it follows that for sufficiently large y, the $\tilde{\kappa}$ term on the LHS can be made arbitrarily small. The LHS can thus be made arbitrarily close to the RHS, so that $g_1(y, p_1)$ becomes arbitrarily close to $g_0(y, p_1)$. Equation 8 gives strong intuition about the validity of the lemma. We provide below a formalized proof of the convergence in the ϵ-δ sense.

We change the arguments of g from how they are defined above, as $p_1 = p$ is constant across g_1 and g_0 and the only element that makes the two different is the presence of $\tilde{\kappa}$. We define the forms that we use here as follows (note that $g(y, \tilde{\kappa}(y)) = g_1(y, p_1)$ and $g(y, 0) = g_0(y, p_1)$ in comparison with the MRS equation).

$$g(y, \tilde{\kappa}(y)) = \text{argmax}_{g\in[0,1]}(1-p)u((1 - \tau + \tau g - \tilde{\kappa})y) + pu((1 - \tau - \tau\theta g - \tilde{\kappa})y) \tag{9}$$

$$g(y, 0) = \text{argmax}_{g\in[0,1]}(1-p)u((1 - \tau + \tau g)y) + pu((1 - \tau - \tau\theta g)y) \tag{10}$$

As κ is a constant, we know that : $lim_{y->\infty}\tilde{\kappa}(y) = 0$.

Then, by the definition of limits, we know that for any $\delta > 0$, there exists a $c \in \mathbf{R}$ such that :

$$y > c \quad \Rightarrow \quad |\tilde{\kappa}(y) - 0| < \delta \tag{11}$$

Set $\epsilon > 0$, by continuity of g on real positive numbers, there exists some $c \in \mathbf{R}$,

$$y > c \quad \Rightarrow \quad |\tilde{\kappa}(y) - 0| < \delta \quad \Rightarrow \quad |g(y, \tilde{\kappa}(y)) - g(y, 0)| < \epsilon \tag{12}$$

Then, $g(y, \tilde{\kappa})$ converges to $g(y, 0)$ as y becomes arbitrarily large. Assumption 1 ensures that, as y becomes arbitrarily large, $g(y, 0)$ will be arbitrarily close to its non-zero limit.

□

Proposition 1. ***High-Income Concealment.*** *Under Assumption 1, there is a cutoff in the model $\hat{y}$ such that holding all else fixed, $y > \hat{y} \implies a = 1$ is optimal.*

Proof. We want to show that for a sufficiently large y, the difference in expected utility between $a = 1$ and $a = 0$ given optimal g_1 and g_0 must be positive. We express expected utility as a function of a and g_a as

$$U(p, \kappa, y) = (1 - p)u((1 - \tau + \tau g(p, \kappa, y) - \tilde{\kappa})y) + pu((1 - \tau - \tau\theta g(p, y, \kappa) - \tilde{\kappa})y), \tag{13}$$

where $g(p, \kappa, y)$ denotes the optimal level of evasion as a fraction of income, e/y, given the primitives. The difference between utility under adoption and non-adoption, given optimal evasion, is simply

$$\Delta_a U = U(p_1, \kappa, y) - U(p_0, 0, y). \tag{14}$$

The key to making use of Lemma 2 is to benchmark these expected utilities to expected utility under $\kappa = 0$ and $p = p_1$ - in which case behavior is given by $g(p_1, y, 0) = g_0(p_1, y)$, and expected utility by $U(p_1, 0, y)$. Adding and subtracting this from both sides of the above expression, we obtain:

$$\Delta_a U = [U(p_1, \kappa, y) - U(p_1, 0, y)] + \{U(p_1, 0, y) - U(p_0, 0, y)\} \tag{15}$$

Equation (15) decomposes $\Delta_a U$ into the difference due to the incursion of the cost - the first term in square brackets - and the difference due to the lower probability of detection - the second term, in curly brackets. The remaining structure of the proof shows that under Assumption 1, the latter dominates the former for large y.

Using the second fundamental theorem of calculus, we can rewrite the term in curly brackets above as

$$U(p_1, 0, y) - U(p_0, 0, y) = -\int_{p_1}^{p_0} U_p(p, 0, y)dp, \tag{16}$$

where $U_p(p,0,y)$ is the partial derivative of U with respect to p evaluated at $(p,0,y)$. Using the envelope theorem to characterize $U_p(p,0,y)$, we have

$$U(p_1,0,y) - U(p_0,0,y) = \int_{p_1}^{p_0} [u((1-\tau+\tau g(p,0,y))y) - u((1-\tau-\tau\theta g(p,0,y))y)]dp. \tag{17}$$

Note that provided $g(p,0,y) \neq 0$ for $p \in [p_1,p_0]$, this expression is strictly positive, because $p_1 < p_0$ and $u' > 0$. In words, provided the individual actually does evade some tax, decreasing the detection probability strictly increases expected utility.

To simplify expressions, as in the proof of Lemma 2, we define the argument of the utility function in the detected and undetected state given behavior $g(p,\kappa,y)$ by $c_D(p,\kappa,y)$ and $c_N(p,\kappa,y)$ respectively. Using equation (17) and the definition of U, we can rewrite equation (15) as

$$\begin{aligned}\Delta_a U =& (1-p_1)[u(c_N(p_1,\kappa,y)) - u(c_N(p_1,0,y))] + p_1[u(c_D(p_1,\kappa,y)) - u(c_D(p_1,0,y))] \\ &+ \int_{p_1}^{p_0} [u(c_N(p,0,y)) - u(c_D(p,0,y))]dp.\end{aligned} \tag{18}$$

Next, we use the second fundamental theorem of calculus again to express all the differences in utilities in the above equation as integrals of marginal utility over the appropriate range of final consumption. To understand these integrals, it helps to note that both $c_N(p,\kappa,y)$ and $c_N(p,\kappa,y)$ are decreasing in κ.[50] We write all integrals so that the lower limit of integration is less than the upper limit.

$$\Delta_a U = -(1-p_1)\int_{c_N(p_1,\kappa,y)}^{c_N(p_1,0,y)} u'(c)dc - p_1\int_{c_D(p_1,\kappa,y)}^{c_D(p_1,0,y)} u'(c)dc + \int_{p_1}^{p_0}\int_{c_D(p,0,y)}^{c_N(p,0,y)} u'(c)dcdp. \tag{19}$$

We now use diminishing marginal utility to find a simpler function $f(y)$ such that $\Delta_a U > f(y)$ always, and then construct an argument that $f(y) > 0$ for sufficiently large values of y. For integrals with a positive sign in front (the third term), we construct f so that the integral is evaluated as a constant at the smallest u' over the specified range, which by $u'' < 0$ corresponds to u' at the upper limit of integration. For integrals with a negative sign in front (the first two terms), we should use the lower limit of integration. We thereby

[50] Differentiating the first-order condition in equation 7, we have $u''(c_N)\frac{\partial c_N}{\partial \kappa} = u''(c_D)\frac{\partial c_D}{\partial \kappa}$. This implies that the sign of $\frac{\partial c_D}{\partial \kappa}$ and $\frac{\partial c_D}{\partial \kappa}$. These two cannot both be positive, because this would imply that evasion is both increasing in κ (from $\frac{\partial c_N}{\partial \kappa} > 0$) and decreasing (from $\frac{\partial c_D}{\partial \kappa} > 0$). Hence they are both negative.

obtain

$$\begin{aligned}\Delta_a U > &-(1-p_1)[c_N(p_1,0,y)-c_N(p_1,\kappa,y)]u'(c_N(p_1,\kappa,y)) \\ &-p_1[c_D(p_1,0,y)-c_D(p_1,\kappa,y)]u'(c_D(p_1,\kappa,y)) \\ &+\int_{p_1}^{p_0}[c_N(p,0,y)-c_D(p,0,y)]u'(c_N(p,0,y))dp. \end{aligned} \tag{20}$$

We modify $f(y)$ slightly by noting that from the first-order condition in equation 7,

$$u'(c_D(p_1,\kappa,y)=u'(c_N(p_1,\kappa,y)\frac{1-p_1}{\theta p_1}.$$

We also note that we can shrink the expression further by evaluating the last term with a constant marginal utility $u'(c_N(p_1,0,y))$, as c_N is decreasing in p and $u''<0$. Substituting this into equation (20) and simplifying, we obtain

$$\begin{aligned}\Delta_a U > &-(1-p_1)u'(c_N(p_1,\kappa,y))\left\{c_N(p_1,0,y)-c_N(p_1,\kappa,y)+\theta^{-1}[c_D(p_1,0,y)-c_D(p_1,\kappa,y)]\right\} \\ &+u'(c_N(p_1,0,y))\int_{p_1}^{p_0}[c_N(p,0,y)-c_D(p,0,y)]dp. \end{aligned} \tag{21}$$

We note that by construction $c_N(p,0,y)-c_D(p,0,y)=\tau(1+\theta)g(p,0,y)y$. As this expression is decreasing in p by Lemma 1, we shrink the function by evaluating it at the upper limit of integration. In so doing we arrive at an $f(y)$ that is simple enough to analyze for large y.:

$$\begin{aligned}\Delta_a U > f(y) \equiv &-(1-p_1)u'(c_N(p_1,\kappa,y))\left\{c_N(p_1,0,y)-c_N(p_1,\kappa,y)+\theta^{-1}[c_D(p_1,0,y)-c_D(p_1,\kappa,y)]\right\} \\ &+u'(c_N(p_1,0,y))(p_0-p_1)\tau(1+\theta)g(p_0,0,y)y. \end{aligned} \tag{22}$$

As $u'>0$ we find that[51]

$$\begin{aligned}f(y)>0 \iff &-(1-p_1)\left\{c_N(p_1,0,y)-c_N(p_1,\kappa,y)+\theta^{-1}[c_D(p_1,0,y)-c_D(p_1,\kappa,y)]\right\} \\ &+\frac{u'(c_N(p_1,0,y)}{u'(c_N(p_1,\kappa,y))}(p_0-p_1)\tau(1+\theta)g(p_0,0,y)y>0 \end{aligned} \tag{23}$$

We now examine the behavior of the expression in equation (23) at large y. We know from Lemma 2 that $c_N(p_1,0,y)-c_N(p_1,\kappa,y)$ and $c_D(p_1,0,y)-c_D(p_1,\kappa,y)$ both become arbitrarily small as y becomes

[51] If u' converges to a strictly positive constant for arbitrarily large y, the proof from this point is more straightforward than what we present here. The result essentially follows directly from Assumption 1 and Lemma 2, which guarantee that the term in the top row shrinks while the term in the bottom row grows large. We construct the proof the way that we do to handle the case where u' approaches zero for large y, which is widely considered to be relevant.

large. The term in the top row can therefore be made arbitrarily small. From Lemma 2, we also know that $\frac{u'(c_N(p_1,0,y))}{u'(c_N(p_1,\kappa,y)}$ converges to unity as y becomes large. Assumption 1 ensures that the second part of the term in the bottom row, $\tau(1+\theta)g(p,0,y)y$ grows arbitrarily large for large y. It follows that $f(y) > 0$ for sufficiently large y, and thus that $\Delta_a U > 0$ for sufficiently large y. □

Proposition 2. ***Incentivizing Concealment.*** *Suppose a policy increases the probability of detection only if $a = 0$. This policy will increase concealment.*

Proof. This result follows immediately from the envelope theorem. Differentiating $\Delta_a U$ with respect to p_0 and applying the envelope theorem, we obtain

$$\frac{\partial \Delta_a U}{\partial p_0} = u(c_N(p_0, 0, y)) - u(c_D(p_0, 0, y)) > 0. \tag{24}$$

□

Proposition 3. ***Comparative Statics of the Resource-Constrained Model.*** *In the optimization problem described by equation (3),*

- $\frac{\partial N_h}{\partial c_h} < 0$
- $\frac{\partial N_l}{\partial c_h} > 0$ if and only if $-N_h R_h'' / R_h' < 1$.

Proof. We solve the resource constraint for N_l in equation(3) and substitute this into the right-hand side of 5. We differentiate the resulting expression with respect to c_h and solve for $\frac{\partial N_h}{\partial c_h}$ to obtain:

$$\frac{\partial N_h}{\partial c_h} = \frac{R_h'(N_h) - N_h R_l''(N_l)}{c_h R''(N_h) - R'(N_h)} < 0 \tag{25}$$

The first result then follows from $R_\theta' > 0$ and $R_\theta'' < 0$ for each type $\theta = 0, 1$.

Proceeding similarly for N_l, we obtain

$$\frac{-(R_h'' N_h + R_h')c_l}{c_l^2 R_h'' + c_h^2 R_l''}. \tag{26}$$

This expression is positive whenever $N_h R_h'' + R_h' > 0 \iff -N_h R_h'' / R_h' > 1$. □

Proposition 4. ***Comparative Statics Without the Resource Constraint.*** *Consider the optimization problem described by equation (3) but ignore the resource constraint. In this model*

- $\frac{\partial N_h}{\partial c_h} < 0$

- $\frac{\partial N_l}{\partial c_h} = 0$.

Proof. These follow directly from differentiating the FOC in equation (6). □