

# The Distribution of Undetected Under-Reported Income Identified by Detection-Controlled Estimation Methods \*

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

November 21, 2023

## Abstract

Official estimates of the tax gap for individual income tax filers in the United States incorporate under-reporting detected in random audits, combined with estimated under-reporting not detected in random audits using a procedure called “Detection Controlled Estimation” (DCE). This paper studies the distributional implications of DCE. We discuss the conceptual challenges in identifying the location of undetected under-reporting through the income distribution. We present results comparing existing methods for distributing undetected under-reporting, and we develop some simple illustrative alternatives. We show how the distribution of undetected under-reporting influences estimates of the rate of under-reporting through the income distribution, the concentration of under-reported income, and the overall concentration of fiscal income.

---

\*Corresponding Author: Daniel Reck, [d.h.reck@lse.ac.uk](mailto:d.h.reck@lse.ac.uk). We thank Drew Johns for helping us to understand and implement the Detection Controlled Estimation methods used in Tax Gap studies. We also thank Will Boning, Daniel Hemel, Janet Holtzblatt, Katie Lim, Steve Rosenthal, David Splinter, Alex Yuskavage, and numerous seminar and conference participants for helpful discussions and feedback. Francesco Armillei 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.

All findings, opinions, and errors are those of the authors alone and do not necessarily represent the opinions of the Internal Revenue Service or Treasury Department.

# 1 Introduction

In [Guyton et al. \(2021\)](#), we estimate the distribution of under-reported income in the United States. The main finding is that the traditional source of data on individual tax filer non-compliance, National Research Program (NRP) risk-stratified random audits, paint an incomplete picture of tax evasion at the top of the income distribution.<sup>1</sup> Evasion via offshore financial assets and evasion in pass-through businesses are relatively rarely detected in these audits, despite evidence they may be substantial. In [Guyton et al. \(2021\)](#) we examine the implications of these findings for income under-reporting and the tax gap through the income distribution.

For this last contribution, about how to modify official estimates of under-reporting and the tax gap, the question arises as to how to incorporate adjustments for under-reporting not detected during audits that are currently included in official statistics (see e.g. [IRS, 2019](#)). In constructing official estimates of aggregate under-reporting and unpaid taxes, IRS researchers estimate a model of auditor effects on random audit data to estimate aggregate undetected under-reporting. These methods were designed to estimate aggregate undetected noncompliance, but not for distributional analysis. How should the resulting aggregate estimates of undetected noncompliance be allocated across the income distribution? Given the concentration of taxable income, it is especially important to understand how much undetected under-reporting may occur by taxpayers in the top 1% of the income distribution, or within the top 1% at higher quantiles.

In the benchmark estimates of the first working paper version of [Guyton et al. \(2021\)](#), released as an NBER working paper in March 2021, undetected under-reporting is distributed following the method of [Johns and Slemrod \(2010\)](#), the first paper to attempt to distribute undetected under-reporting through the income distribution.<sup>2</sup> This particular method of distributing undetected under-reporting has been the subject of some recent criticism ([Hemel et al., 2021](#); [Auten and Splinter, 2021](#)). Our approach in the initial draft of [Guyton et al. \(2021\)](#) was the best approach given the data and methods available to us at the time, but the criticism rightly identifies an important gap in our understanding of DCE methodology. These DCE methods were designed to identify under-reporting and the tax gap in the full population, while identifying undetected under-reporting at the distributional level requires further assumptions. As such, applying existing methods uncritically to distributional analysis risks making implicit assumptions that may be unrealistic. This issue becomes particularly important when we focus on the very top tail of the income distribution, where relatively few individuals report a large share of income ([Piketty and Saez, 2003](#)).

---

<sup>1</sup>We stress that it is not the randomness of these audits that is the issue. Random sampling is a fundamental tool for tax compliance research. The issue we highlight in our paper is that audit procedures and the information and resources available to auditors limit the detection of some forms of non-compliance.

<sup>2</sup>[DeBacker et al. \(2020\)](#) implement the same method as [Johns and Slemrod \(2010\)](#) on similar random audit data to ours, with similar results. To our knowledge, these are the only two other papers to include such undetected under-reporting in distributional analysis of tax compliance.

In this technical note, we revisit the issue of the distribution of undetected under-reported income. We aim to make progress on some immediate questions about how to allocate undetected under-reporting identified by DCE through the income distribution, toward a revision of [Guyton et al. \(2021\)](#). We also hope to motivate further analysis of this question by identifying the deeper challenges involved with identifying undetected under-reporting at the distributional level. We discuss the underlying conceptual issues around identifying undetected under-reporting at the distributional level, and we comprehensively analyze several alternative methods. We begin with a review of existing methods for identifying and distributing undetected non-compliance using “Detection Controlled Estimation” (DCE), in official statistics and prior literature. To assist our understanding of the strengths and weakness of use of existing DCE methods for distributional analysis, we show how existing methods shape the profile of under-reporting at the distributional level, and we present them alongside some alternative methods, in which we make explicit distributional assumptions for illustrative purposes. We conclude with a discussion of some further methodological issues that we do not resolve here but flag for future research. Throughout, we set aside the main question from [Guyton et al. \(2021\)](#)—namely how accounting for sophisticated forms of evasion via offshore accounts and pass-through businesses modifies official estimates of non-compliance. Rather, we focus on identifying the location in the income distribution of undetected under-reporting that *is* included in official statistics.

The core exercise we undertake here is to compare two existing “micro” approaches for accounting for undetected under-reported income with some illustrative “macro” approaches to the same question. We implement two existing micro approaches on the same random audit data, from tax years 2008–2013. Specifically, we implement the approach used in [IRS \(2007\)](#), originally implemented using tax year 2001 random audit data (“DCE2001”), and the one used in [IRS \(2019\)](#), originally implemented using 2008–2013 random audit data (“DCE2019”). By allocating undetected under-reporting of each type of income at the micro (individual tax return) level, one can straightforwardly map under-reported income to the tax gap, which is the main purpose of these methods in Tax Gap studies. For analysis of under-reporting throughout the income distribution, however, the structure imposed by either micro approach on the relationship between the probability of detection of under-reporting and true income becomes crucial. It can be difficult to assess what structure is implicitly imposed on this relationship by micro allocation methods. As such, we compare the results from micro approaches with macro approaches, in which we make explicit and direct assumptions about the relationship between undetected under-reporting and true income (at the cost of making the map between under-reported income and the tax gap more opaque). For these macro approaches, we take the total undetected under-reporting of each type of income as given and distribute this under-reporting to different parts of the income distribution, assuming that undetected under-reporting exhibits a similar concentration to income distributions that we do observe. For example, we might assume that undetected

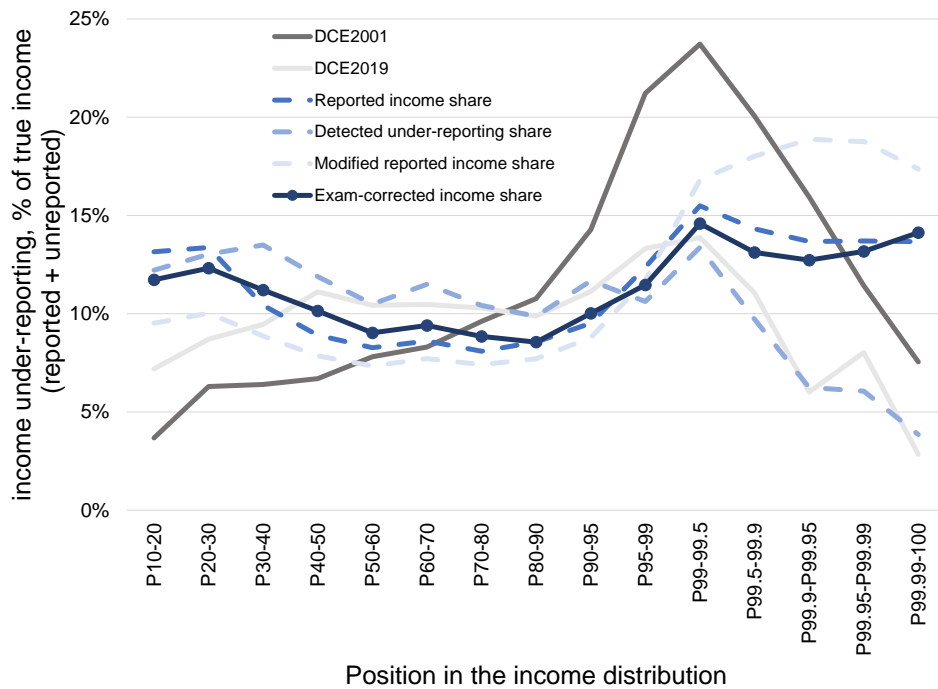
under-reporting of a given type of income is distributed like reported amounts of the same type of income – e.g. if X% of interest income belongs to the top 1% of the reported income distribution, we would allocate X% of undetected under-reported interest income to the top 1%. We also illustrate the implications of assuming that undetected under-reported income is distributed like exam-corrected incomes (incomes after adjusting for detected under-reporting from random audits), like exam-detected under-reporting (under-reporting detected in random audits), or like a modification of reported income in which we regard some similar types of income as a single income class (see details below).

The main takeaways from our analysis are summarized in Figure 1, which we unpack in depth below. In this figure, we compare under-reported income as a share of true income, ranking individuals by estimated true income, for these micro and macro allocation mechanisms. Doing so reveals the structure imposed by micro approaches on the relationship between detection and true income. Specifically, we find that the approach used in the most recent Tax Gap study (IRS, 2019), implicitly imposes that undetected under-reporting exhibits the same concentration as exam-detected under-reporting. Because there is very little undetected under-reporting at the top, supposing that undetected and detected under-reporting are similarly distributed results in relatively little undetected under-reporting being allocated to the top. Meanwhile, the older approach employed in IRS (2007), and in Johns and Slemrod (2010), appears to assume that undetected under-reporting is distributed like reported incomes, but compared to the macro allocation method based on the concentration of reported incomes, the profile of under-reporting with this micro approach is more steeply increasing with income in the bottom 99% and more steeply decreasing with income in the top 1%. In what follows, we attempt to understand how the micro approaches implicitly impose such structure on the relationship between detection probabilities and true income, what theory and other data suggest about this relationship, and the potential limitations of both micro and macro approaches.

## 2 Overview and Prior Work on DCE

The method used in IRS Tax Gap studies to account for undetected under-reporting of income (and the resulting unpaid tax) is known as “Detection Controlled Estimation,” or DCE. Starting from data on *detected* non-compliance from stratified random audits of a representative sample of tax returns, the aim is to identify under-reporting that was present but not detected during these audits, i.e., *undetected* under-reporting. The core principle of DCE is that the probability of detection via audit of a given dollar of under-reporting is typically less than one. The objective of the method is to extrapolate to total non-compliance in a counterfactual scenario where the probability of detection equals one for every dollar of under-reporting. Obtaining empirical traction on this extrapolation requires exogenous variation in detection, e.g., some instrumental

**FIGURE 1: COMPARING MICRO AND MACRO ALLOCATION METHODS: DISTRIBUTIONAL ESTIMATES FOR INCOME UNDER-REPORTING AS A % OF TRUE INCOME, TY2008–2013**



Note: This figure illustrates the main comparison of different allocation methods for undetected under-reporting. We compare the micro allocation methods used in prior Tax Gap studies (DCE 2001 and DCE2019) with macro allocation methods supposing that undetected under-reporting is distributed like some other type of income we observe (see main text for details). We observe that DCE2019 implicitly imposes that undetected under-reporting exhibits the same concentration as exam-detected under-reporting. Meanwhile, DCE2001 imposes that undetected under-reporting is distributed like reported incomes, but compared to the macro allocation method based on the concentration of reported incomes, the profile of under-reporting with this micro approach is more steeply increasing with income in the bottom 99% and more steeply decreasing with income in the top 1%. Figures 4, 5, and 7 below and their descriptions contain further illustrations and details on these various methods.

variable that affects the likelihood of detection and is otherwise not associated with under-reporting. Current methods obtain this via the identity and characteristics of the auditor, estimating therefore a structural model of differential detection by the auditors conducting the audits.

One assumption behind current DCE methods is that the best auditors detect all under-reporting, i.e., that there is a mass of auditors with 100% detection probabilities for every dollar of under-reporting on any return. The methods extrapolate detected under-reporting to a counterfactual in which all auditors are replaced by such auditors. The evidence in [Guyton et al. \(2021\)](#) (henceforth GLRRZ) suggests that this assumption might be too strong, at least at the top of the income distribution. Nevertheless, if we were to relax the assumption to account for the possibility that some under-reporting is undetectable by virtually any auditor, DCE methods would still identify *detectable but undetected* under-reporting, and it remains important to assess how large this under-reporting is and where it locates in the income distribution.<sup>3</sup>

A second assumption is that auditor assignment is orthogonal to the determinants of under-reporting and detection, in other words that auditors are effectively randomly assigned. Under this assumption, systematic between-auditor variation in detected under-reporting is attributable to differential detection rates, and not simply the selection process assigning auditors to cases. A concern with this assumption is that auditors may in fact be assigned in a selected fashion, which we discuss in Section 5.

The DCE methods used in official statistics and related research have evolved over time.<sup>4</sup> [Feinstein \(1991\)](#) describes the first auditor-effects DCE method. Official tax gap statistics in [IRS \(2007\)](#) employ such an auditor effects model with data from the tax year 2001 wave of NRP random audits. These same data and methods are used in [Johns and Slemrod \(2010\)](#), who discuss the key details of this method and apply it to distributional questions. The procedure for this version of DCE is 1) to estimate multipliers representing the ratio of total under-reporting to total detected under-reporting within a class of taxpayers and types of income, and 2) to scale detected under-reporting by such multipliers. We refer to this method going forward as the *DCE2001* method, because it is primarily associated with the 2001 NRP data. The March 2021 draft of GLRRZ and [DeBacker et al. \(2020\)](#) implement DCE2001 methods on 2006–2013 NRP data.<sup>5</sup>

Newer DCE methods aim to make improvements on DCE2001, with a focus in particular on the “extensive margin” of under-reporting. For a given type of income (e.g. wages), many individuals have zero detected under-reporting. However, some of these individuals might have undetected under-reporting. Scaling detected under-reporting at the micro level to account for undetected under-reporting, as in DCE2001,

---

<sup>3</sup>A related question is how to account for potential double counting between undetected under-reporting implied by DCE and the sophisticated evasion considered by GLRRZ. We defer this question to GLRRZ, which discusses adjustments to avoid double counting.

<sup>4</sup>A complete set of tax gap reports and descriptions of underlying methods from 1988 to present is available at <https://www.irs.gov/statistics/irs-the-tax-gap>.

<sup>5</sup>A caveat with this approach is that the data-generating process may have changed from 2001 to 2006–2013 in such a way that the older methods tailored to the 2001 wave of NRP may no longer give an accurate estimate of undetected under-reporting when applied to later years. We should keep this in mind when comparing results employing DCE2001 to other methods below.

does not capture such extensive margin variation in under-reporting or detection. This issue does not necessarily introduce bias in estimates of total income under-reporting, but given the non-linear income tax schedule, this issue matters for mapping income under-reporting to the tax gap. To address this issue, the two most recent tax gap studies, using waves of NRP data from tax years 2008–2010 and 2011–2013, estimate by maximum likelihood a structural model which incorporates an extensive margin of both detection and under-reporting (IRS, 2016, 2019). The method used in these recent tax gap studies is described in Erard and Feinstein (2011).<sup>6</sup> In this paper, we build on these revised DCE methods, building on the estimates in IRS (2019). We refer to these newer methods as the *DCE2019* method.

We obtained the simulated data underlying the individual income tax gap estimates for tax years 2011–2013 (IRS, 2019, see Section 4.2 and especially Table 5). These simulations used 2008–2013 NRP data to estimate the tax gap for tax years 2011–2013. Our analysis in GLRRZ includes data from 2006–2013; we were unable to extend DCE2019 to the tax year 2006–2007 wave of the NRP. The aggregate figures we report here are population-weighted averages over tax years 2008–2013.<sup>7</sup> When distributing undetected under-reporting in the next section, we continue to focus mainly on tax years 2008–2013, but we also construct some estimates of distributional statistics to represent the full sample period in GLRRZ, tax years 2006–2013, assuming that the rates of under-reporting (as a share of true income) for different types of income were stable over time in this period.

### 3 Estimates of Total Detected and Undetected Under-Reporting

We begin by illustrating how adding estimated undetected under-reporting modifies the size and composition of income under-reporting and the tax gap, using aggregate under-reporting according to the DCE2019 method. We turn to the question of how under-reporting is distributed through the income distribution in the next section.

In total, an estimated 10.7% of estimated true income is under-reported.<sup>8</sup> Exam-detected under-reporting comprises 3.8% of estimated true income, and estimated undetected under-reporting comprises the remaining 6.9%. In 2012 dollars, 975\$ billion of income was estimated to be under-reported and 300\$ billion of individual income tax (including self-employment taxes and refundable credits) unpaid on average over tax years 2008–2013. The latter figure closely matches the official tax gap figure for 2011–2013 in IRS (2019), \$290 billion.<sup>9</sup> This is unsurprising because these results are built on the same data and programs as the 2019

<sup>6</sup>There was also an intermediate version of DCE methods employed in estimates of the tax gap circa 2012, using NRP data from tax years 2001 and 2006. See Bloomquist et al. (2012) for details.

<sup>7</sup>The fact that we construct estimates to represent 2008–2013 while the official statistics constructed estimates to represent 2011–2013 makes a negligible quantitative difference to aggregate statistics of interest.

<sup>8</sup>This and all other figures for under-reporting in this paper are net of over-reporting.

<sup>9</sup>See IRS (2019, Table 2 p. 11): the individual income underreporting tax gap is \$245 billion and the self-employment underreporting

Tax Gap study.

Figure 2 breaks down the aggregate adjustment by type of income, showing how DCE adjustments modify estimated under-reporting of each type of income relative to estimates before DCE adjustment. To facilitate interpretation, we scale the dollar totals by estimated total true income, including detected and undetected under-reported income. Overall, as already noted, the DCE adjustments (sum of the blue bars) add up to 6.9% of estimated true income. The aggregate adjustment is significantly larger than the amount of exam-detected noncompliance (3.8% of true income, sum of the grey bars). We can see on Figure 2 that the bulk of the aggregate DCE adjustment comes from adjustments to Schedule C income, capital gains, rental income, pass-through business income, and line 21 other income (e.g. net operating loss carry-forwards), forms of income that are moderately to highly concentrated at the top of the distribution.<sup>10</sup> The DCE adjustments for these income categories combined account for more than 80% of the aggregate DCE adjustment. For pension income and wage income (forms of income that are relatively equally distributed<sup>11</sup>), the adjustment is relatively small (less than 10% of the aggregate DCE adjustment). The DCE adjustment significantly increases estimated noncompliance for forms of income that are subject to incomplete or no third-party information reporting (such as business income); less so for forms of income subject to significant third-party information reporting (such as wages and pension income).

The composition of overall income and total under-reported income changes slightly when we incorporate undetected under-reporting. Table 1 summarizes the composition of income according to exam-corrected NRP data without accounting for undetected under-reporting for 1) the set of years in GLRRZ, 2006–2013, and 2) the set of years for which we have DCE2019 estimates, 2008–2013. There are some minor differences across years in the composition of overall and under-reported income. Table 2 reports analogous estimates including undetected under-reporting according to DCE2019. To understand how DCE changes the composition of overall and under-reported income, we should compare the estimates in Table 2 to those from Table 1 for 2008–2013. Doing so, we observe that the biggest change in the composition of overall income due to incorporating undetected under-reporting are that schedule C income increases in importance from 5.3% of estimated total income to 7.6%, while wages decrease in importance from 72.4 to 66.3% of estimated total income.<sup>12</sup> We also observe modest increases in the importance of capital gains, Form 1040 Line 21 “Other Income” (which includes e.g. NOL carryforwards), and rental income as a share of total income. The composition of *under-reported* income changes in a somewhat different fashion: Schedule C

---

tax gap an additional \$45 billion.

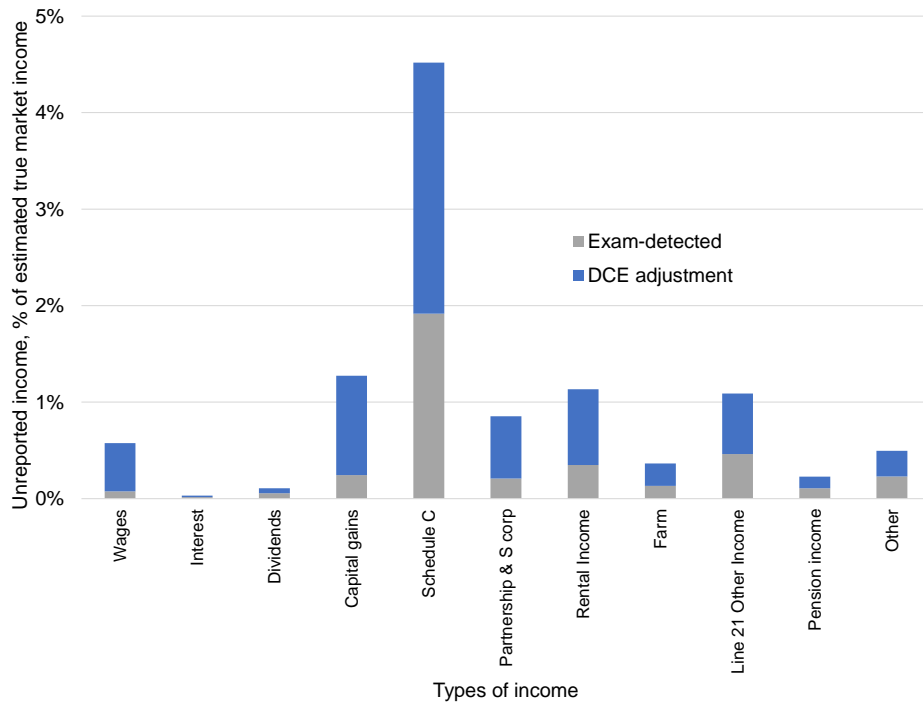
<sup>10</sup>In 2012, the top 1% of tax filers earned 22.7% of reported market income, 19% of schedule C income, 83% of capital gains, 78% of pass-through business income, and more than 100% of rental income (which on aggregate was negative).

<sup>11</sup>In 2012, the top 1% of tax filers earned 12% of wage income and 6% of taxable pension income excluding Social Security.

<sup>12</sup>These figures are similar to but not directly relatable to those appearing in Figure 2 because the estimates in Figure 2 scale both detected and undetected under-reporting by estimated total true income including undetected under-reporting, while the figures in the first column of Table 1 scale by total exam-corrected income, not including undetected under-reporting.



**FIGURE 2: UNDER-REPORTED INCOME BY TYPE OF INCOME, % OF ESTIMATED TRUE MARKET INCOME (DCE2019)**



Note: This figure illustrates estimated detected and undetected under-reporting in DCE2019 by type of income. We scale both exam-detected and DCE adjusted under-reporting totals by estimated total true income (= reported income + detected under-reporting + undetected under-reporting). Note also that what [IRS \(2019\)](#) call “Form 1040 Other Income” 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 last bar of the figure refers to all other components of income.

income declines in importance from 50.5% of detected under-reporting to 42% of detected plus undetected under-reporting, while capital gains, partnership and S corp business income, and wages increase increase in relative importance after accounting for undetected under-reporting.

We note changes in methods from DCE2001 to DCE2019, and possibly changes in underlying data, contribute to some differences in estimated total under-reporting across different versions of DCE methods. Based on comparisons of our work implementing older DCE methods (see the first working paper version of GLRRZ, Table A2), we observe that the newer methods for DCE adjustment decrease overall estimated under-reporting as a share of true income slightly, from about 14% to 11% of true income under-reported. Breaking things down by type of income, we observe that compared to DCE2001, the new methods in DCE2019 increase estimated rates of under-reporting for capital gains (from 13% to 23%) and decrease estimated rates of under-reporting for sole proprietor income (from 68% to 60%) and partnership and S corp business income (from 22% to 15%). All of this is consistent with the descriptions of methods and results reported in official Tax Gap studies using DCE2001 and DCE2019. The biggest differences between the methods appear to be mainly attributable to the change in the classification of types of income used in the different versions of DCE, see page 3 of [IRS \(2007\)](#) and Table 4 of [IRS \(2019\)](#). For example, capital gains was classified as "high-visibility" income in DCE2001 and grouped with other forms of income subject to extensive third-party information-reporting in the estimation of the model, but capital gains were grouped with forms of income *not* subject to third-party reporting in the newer DCE2019. The rationale for this particular change is discussed in [Bloomquist et al. \(2012\)](#).

## 4 Distributing Undetected Under-Reporting: Methods and Estimates

Next, we take these aggregate figures on the extent of under-reporting of each type of income from the previous section as given, and turn to the question of how undetected under-reporting is distributed through the true income distribution. Here, we confront in our view the central challenge in incorporating the undetected evasion from official statistics into distributional estimates of income under-reporting and the tax gap. Existing work leaves room for uncertainty about where in the income distribution undetected under-reporting belongs, especially when it comes to the very top of the distribution. Moreover, as we have just seen, estimated under-reporting is large in the aggregate, so its location exerts significant influence on the overall distribution of the tax gap. In this section we will present and critically assess several alternative methods for estimating undetected evasion at the distributional level.

TABLE 1: TAX EVASION DETECTED IN NRP RANDOM AUDITS WITHOUT DCE CORRECTION: DECOMPOSITION BY INCOME TYPE

	2006-2013 Data				2008-2013 Data			
	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 under-reported income (%)	Total under-reported income of this type/ Total under-reported income (%)	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 under-reported income (%)	Total under-reported income of this type/ Total under-reported income (%)
Capital gains	5.8	7.1	0.28	4.8	4.7	6.4	0.3	5.5
Dividends	3.9	2.8	0.11	2.8	2.2	1.5	0.1	2.8
Interest	1.9	0.7	0.03	1.5	1.6	0.4	0.0	1.1
Line 21 Other income	0.2	11.9	0.47	253.6	0.3	12.1	0.5	195.6
Partnership and S Corp	5.6	6.5	0.26	4.6	5.6	6.8	0.3	5.0
Rental	0.7	9.0	0.35	48.3	0.8	9.2	0.4	48.2
Schedule C	5.3	49.4	1.95	36.8	5.3	50.5	2.1	38.7
Wages	72.4	3.5	0.14	0.2	70.6	2.0	0.1	0.1
Other	4.1	9.2	0.37	8.8	9.0	11.1	0.45	5.0
Total	100.0	100.0	3.96	100.0	100.0	100.0	4.08	100.0

Note: This table describes the composition of overall and under-reported income based on exam-corrections only, before any correction for undetected noncompliance. We report results separately for the set of years used in GLRRZ, tax years 2006–2013 (left panel) and the set of years for which we have DCE2019 estimates, 2008–2013 (right panel). The first column reports the composition of income, including detected under-reported income. The composition of overall income by type of income in the first column is similar to income shares we observe in SOI data, but these shares include detected under-reporting here. Consequently the largest differences with SOI income shares are observed for types of income with significant detected evasion. The second column describes the composition of *under-reported* income, the third scales total under-reported income of a given type by total estimated true income, and the final column reports a type-specific under-reporting rate by scaling total under-reporting of a given type by total true income of that type of income. Note also that what IRS (2019) call “Form 1040 Other Income” 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. We note that the estimated rate of under-reporting by type of income in the fourth column exceeds 100% for Line 21 income. This occurs because Line 21 income can be negative; large negative values are common at the bottom of the income distribution because of net operating losses from pass-through businesses carried through to other years. Large corrections to line 21 are typically disallowed NOLs from other years.

**TABLE 2: THE COMPOSITION OF REPORTED AND UNREPORTED INCOME INCLUDING UNDETECTED UNDER REPORTING USING DCE2019 METHODS, 2008-2013**

	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	5.4	11.9	1.27	23.4
Dividends	2.1	1.0	0.11	5.2
Interest	1.5	0.3	0.03	2.1
Line 21 Other Income	0.9	10.2	1.09	126.1
Partnerships and S Corp	5.6	8.0	0.85	15.3
Rental	1.5	10.6	1.13	75.2
Schedule C	7.6	42.3	4.52	59.8
Wages	66.3	5.4	0.58	0.9
Other	9.3	10.2	1.09	11.8
Total	100.0	100.0	10.7	

Note: This table describes the composition of overall and under-reported income, including undetected non-compliance as estimated by DCE2019, for tax years 2008–2013. The first column reports the composition of income, including detected under-reported income. The second column describes the composition of *under-reported* income, the third scales total under-reported income of a given type by total estimated true income, and the final column reports a type-specific under-reporting rate by scaling total under-reporting of a given type by total true income of that type of income. The under-reporting rates by type of income in the fourth column are very similar to the “Net Mis-reporting Percentages” in Table 5 of IRS (2019). The differences are entirely attributable to a slight difference in definitions: we scale under-reporting by total net income for the given line item, while the Tax Gap statistics scale under-reporting by the total of the absolute value of the given line item. This makes a very minor difference in the mis-reporting rate for all line items except those where income is often negative, such as Rental Income and Line 21 Other Income (see also the note to the previous table).

#### 4.1 Defining The Problem

The fundamental difficulty here is that we observe detected under-reporting at the micro, tax-return level, but we do not observe undetected under-reporting at the micro level. Conceptually, undetected under-reporting belongs in the part of the (true) income distribution where the probability of detection is lower. Keeping this principle in mind helps us evaluate any method for allocating undetected under-reporting through the income distribution. What does the method imply is the relationship between detection probabilities and true income? The answer to this question shapes the distributional estimates generated by the method.

What should we expect the relationship between detection probabilities and true income to be like? Theory and the available data do not provide a definitive answer to this question, but they do provide some useful guidance. As we discussed in the theoretical section of GLRRZ, in the workhorse model of tax evasion by Allingham and Sandmo (1972), risk preferences shape how individuals’ appetites for the (financial)

risks involved with tax evasion depend on their true income. Tax evasion is modeled as financial risk-taking and relative risk aversion governs what fraction of their true income individuals are willing to put at risk. Typically, we would expect relative risk aversion to remain constant or decrease as we move toward the top of the income distribution, which suggests that the rate of under-reporting should be constant or increasing with true income (see Lemma 1 of GLRRZ). Random audit estimates including only detected under-reporting tell the opposite story, especially in the top 1% of the distribution (see Figure 1 of GLRRZ), so the theoretical argument based on relative risk aversion would suggest that overall under-reporting should be more concentrated at the top than exam-detected under-reporting.

This argument based on classical theory is obviously not definitive, but there are additional reasons to suspect that undetected under-reporting should be more concentrated than detected under-reporting. For example, we expect the sophistication of under-reporting to increase with income, based on the main theoretical argument in GLRRZ – the supply-side model of sophisticated tax evasion in [Alstadsaeter et al. \(2019\)](#) also makes this prediction. If sophistication increases with income and detection probabilities decrease with sophistication, this will tend to increase the concentration of undetected under-reporting at the top relative to detected under-reporting. Relatedly, [Phillips \(2014\)](#) presents a model in which simple forms of under-reporting (e.g. not reporting a source of self-employment income at all) are more common at the bottom of the income distribution, while more nuanced forms of under-reporting should be more common further up in the distribution due to heightened scrutiny from the tax authority, and he shows that this is consistent with evidence from random audit data.

Another line of reasoning concerns the composition of income. Existing evidence suggests that the rate of non-compliance is strongly affected by the presence of third party information ([Kleven et al., 2011](#); [IRS, 2016, 2019](#)). Figure 3 decomposes the composition of income into categories based on the extent of third-party information reporting on each type of income, using the same four categories as prior Tax Gap studies. The scope for non-compliance in the bottom 90% of the income distribution comes virtually entirely from Schedule C Sole Proprietor income. This income is supported by “little to no” third party information reporting, so there is substantial scope for non-compliance. But ultimately this type of income only reflects 10% to 15% of income throughout the bottom 90%. The remainder is predominantly wage and salary income with minimal scope for misreporting due to the presence of “substantial information and withholding.” At the top of the distribution, wage and salary income declines substantially in importance, especially within the top 1%, and pass-through business income and capital gains become much more important. Both of these types of income are supported by limited third party information reporting; prior Tax Gap studies call put these in the “some information” category, because there is some third-party information reporting that is useful in enforcing these taxes (e.g. Forms 1099-B for some forms of capital gains, Schedules K-1 for pass-

through businesses), but third-party information is generally not enough to completely cross-validate the income reported by the taxpayer with information from third parties. A Schedule K-1 from a pass-through business, for example, is not clearly from a “third party” if the taxpayer is in a position to control what the business they own reports, as many very high-income individuals are. As a result of all this, a much smaller share of income at the very top of the distribution is supported by comprehensive third party information reporting, relative to the bottom 95% of the distribution.<sup>13</sup> For closely related reasons, we observed in the previous section that estimated undetected under-reporting is concentrated among types of income that are more concentrated at the top of the distribution.

One caveat that could point in the opposite direction, toward less concentrated undetected under-reporting at the top of the distribution, concerns the wider possibilities for tax avoidance at the top. For example, in an extreme scenario where people with very high economic incomes could *always* legally avoid large tax bills on said incomes, they might prefer the legal avoidance route to non-compliance. While it is true that high-income people often have more opportunities for tax avoidance, there is no evidence to our knowledge that legal avoidance crowds out non-compliance. An equally plausible scenario is that after exhausting all legal avoidance opportunities, high-income taxpayers still facing a large tax bill would turn to riskier and less compliant means of sheltering income from taxes. Given that about 36% of all taxes paid come from individuals in the top 1% of the income distribution, clearly some high-income individuals face large tax bills that they do not legally avoid. In any case, we acknowledge the caveat and hope that future research will analyze this further.

Having thought through the question from a conceptual point of view, we now turn to actual distributional estimates using various methods. Different methods for allocating undetected under-reporting make different assumptions about the relationship between true income and detection probabilities. As we review different methods, we focus our discussion on the structure they impose, implicitly or explicitly, on this relationship. We will present results for two categories of methods to allocate a given amount of total under-reporting through the income distribution. We refer to these as micro allocations and macro allocations. We review and present results for each of them in turn.

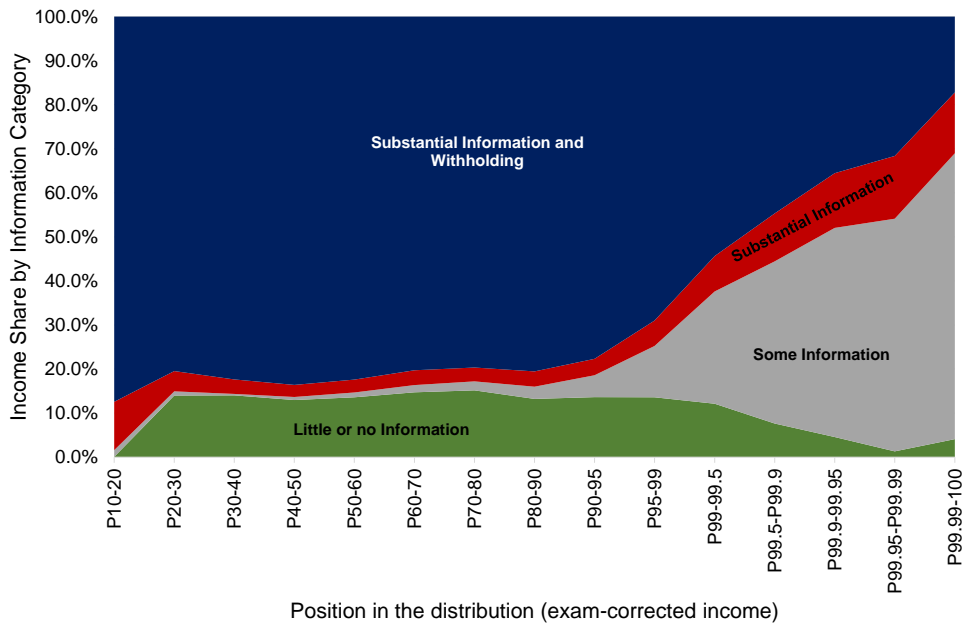
## 4.2 Micro Allocations

The first approach is to use a model to directly assign under-reporting to individuals at the micro level in a representative sample, and then to estimate population-level distributional statistics. We label this a

---

<sup>13</sup>One additional fact reinforces the point we make here. Interest and dividend income are classified as having “substantial information reporting” in Figure 3 and prior studies, due to the reporting of these types of income on 1099 forms by financial institutions. But an increasing share of interest and dividend income within the top 1% of the distribution is not covered by 1099 forms (e.g. income from offshore accounts), which increases the scope for non-compliance. The same is true of 1099-B reporting for capital gains.

**FIGURE 3: THIRD-PARTY REPORTING AND THE COMPOSITION OF INCOME**



Note: This figure divides market income into the four categories used by IRS (2019) and prior Tax Gap studies to describe the comprehensiveness of third-party information and withholding for each type of income. “Little or no information” includes non-farm proprietor income (Schedule C), line 21 other income, rents and royalties, farm income, and Form 4797 income. “Substantial information” includes interest and dividends. “Some information” includes partnership and S corp business income and capital gains. “Substantial information and withholding” includes just wage and salary income. Figure 3 of IRS (2019) uses the same categories and shows estimated under-reporting rates of 55% for “little or no information,” 17% for “some information,” 5% for “substantial information,” and 1% for “substantial information and withholding.” The Figure is based on NRP random audit data for 2006–2013, ranking individuals by exam-corrected income. Ranking by reported or DCE-adjusted income (using DCE2001 or DCE2019) yields very similar results. We observe in the figure a shift in the composition of income from the bottom 99% to the top of the top 1%, away from income with substantial information reporting and toward income with only limited information reporting (e.g. Schedule K-1 reporting for partnership sand S-corps, 1099-B reporting for some but not all capital gains).

*micro*, or bottom-up approach. Both DCE2001 and DCE2019 use a micro approach to allocate undetected under-reporting to specific individuals, mainly to map total income under-reporting to total unpaid tax. The method by which they do so builds on their underlying structural DCE model.

For any micro allocation, the crucial relationship between detection probabilities and true income is *implicitly* determined by the structure the underlying structural model and/or the allocation method that builds on this model imposes on this relationship. Existing micro approaches allocate un-reported income using information from originally filed tax returns, and sometimes information from exam-corrected returns.<sup>14</sup> Because individuals should be ranked by their true income in estimating distributional statistics, the extent of *re-ranking* implied by the allocations of unreported income becomes crucial. Individuals with larger amounts of mis-reporting should be ranked higher in the true income distribution than they are in the reported or exam-corrected income distribution. Micro approaches should obviously incorporate some re-ranking because reported incomes are mechanically selected on compliance, but just how much re-ranking occurs in reality is crucial, and, unfortunately, not well-identified for undetected under-reporting.

To understand how the implicit structure of a micro allocation influences the estimated profile of under-reporting, one informative exercise is to compare results based on DCE2001 and DCE2019. How do these approaches differently shape the location through the distribution of undetected evasion? What does this imply about the structure they impose on the relationship between detection and true income?

Johns and Slemrod (2010) describe the allocation method used in DCE2001, and show how implementing it influences the distribution of unreported and true income. In this allocation method, each initial dollar of exam-detected under-reported income is scaled by a multiplier for the applicable class of taxpayer and type of income. The underlying model identifies the DCE multipliers, and the allocation method assigns the scaled total under-reporting for each initial dollar of detected under-reporting to the same individual who under-reported the initial dollar.

We implement this approach in Figure 4, obtaining a similar estimated profile of evasion to Johns and Slemrod (2010) (see panel (b)). The estimates incorporate substantial re-ranking, because individuals with large amounts of detected under-reporting are allocated large amounts of undetected under-reporting. We unpack the impact of re-ranking in panel (c) of Figure 4. The DCE2001 multiplier approach presumes that individuals with large amounts of exam-detected under-reporting generally have more undetected under-reporting as well, i.e. that exam-detected under-reporting is a strong signal for the presence of undetected under-reporting. This assumption shapes the resulting profile of evasion. Namely, the bulk of exam-

---

<sup>14</sup>None of the existing micro approaches use the identity or characteristics of the auditor conducting the audit to assist in allocating unreported income. Using this information to assist in distributional estimation would be more difficult than one might expect, because the DCE auditor effects model was not initially estimated on the full sample. However, we note that this information might be useful in future attempts at micro allocation of DCE adjustments. This particular issue is discussed further in Reck et al. (2021).



detected under-reporting then falls in the p50-p95 part of the exam-corrected income distribution. Many individuals with exam-corrected incomes in this range are thus allocated significant undetected under-reporting by the method. Re-ranking then pushes much of this under-reporting to the 95-p99.5 part of the distribution. The resulting profile of evasion from the DCE2001 approach increases rapidly up to p99 and then sharply decreases within the top 1% of the distribution – insufficient re-ranking at the very top occurs to reverse the sharp decline in the profile of evasion in estimates based on exam corrections alone.

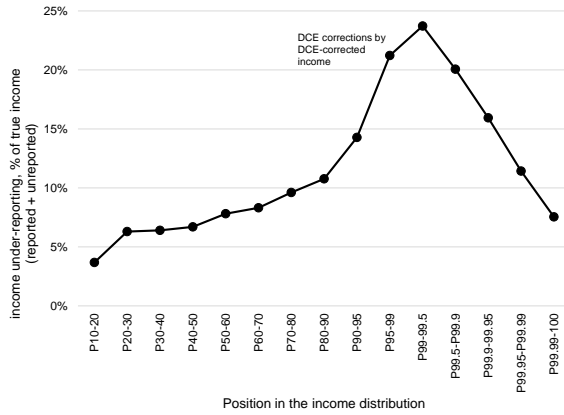
Taken literally, the allocation method used in DCE2001, assuming undetected under-reporting scales linearly at the individual line-item level with detected under-reporting, is plainly unrealistic, as noted by [Johns and Slemrod \(2010\)](#); [DeBacker et al. \(2020\)](#) and [Auten and Splinter \(2021\)](#). Surely some individuals will be allocated too much under-reporting and others too little. One particular unrealistic feature of this allocation is that there is no allowance for an extensive margin of detection: the model does not assign undetected under-reporting to any individuals with zero exam-detected under-reporting. While when taken literally, this allocation is unrealistic at the micro level, the extent and direction of the bias in the distributional aggregate estimates implied by the allocation method is ambiguous. For an extended discussion of this issue and some additional sensitivity analysis on DCE2001, we refer readers to [Auten and Splinter \(2021\)](#) and [Reck et al. \(2021\)](#). The direction of the bias may not even be monotonically related to income: based on Figure 1, it appears that DCE2001 allocates much more undetected under-reporting to the P90-P99 part of the income distribution than we might expect given the distribution of reported incomes or exam-corrected incomes, as we discuss further below.

Newer work in DCE2019 employs a different method, leading to a different relationship between detection probabilities and true income. Estimates of the tax gap in official statistics were based on simulations that in turn built on the underlying model of DCE, as alluded to on page 18 of [IRS \(2019\)](#). These simulations are described as an improvement to previous methods because they incorporate an extensive margin adjustment, which would imply a better mapping of income-under-reporting and the tax gap. We obtained the programs for these simulations and implemented them, to further see what the micro-simulated data would imply about our distributional statistics of interest.

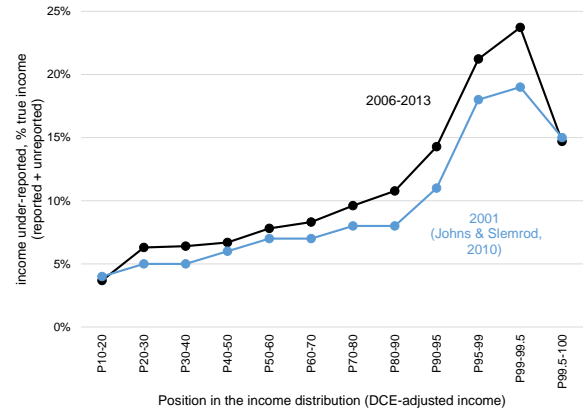
The DCE2019 micro allocation does not use information from micro-level exam corrections to allocate DCE adjustments. Rather, it relies on simulations built on the estimated structural model of DCE. As discussed above, the extensive margin of under-reporting, i.e. whether there is any under-reporting present on a given line, plays a key role in this model of DCE. Before the simulation, based on observables  $X$  on the tax return each income line item on each individual tax return is allocated 1) a probability of containing under-reported income,  $P(\Delta y > 0|X)$ , and 2) a dollar amount of “expected” under-reporting that will be allocated to them if they did in fact under-report,  $\hat{y}_{cor}(X)$ . Both of these quantities are based on the estimated

**FIGURE 4: DISTRIBUTIONAL ESTIMATES FOR INCOME UNDER-REPORTING AS A % OF TRUE INCOME USING DCE2001, TY2006–2013**

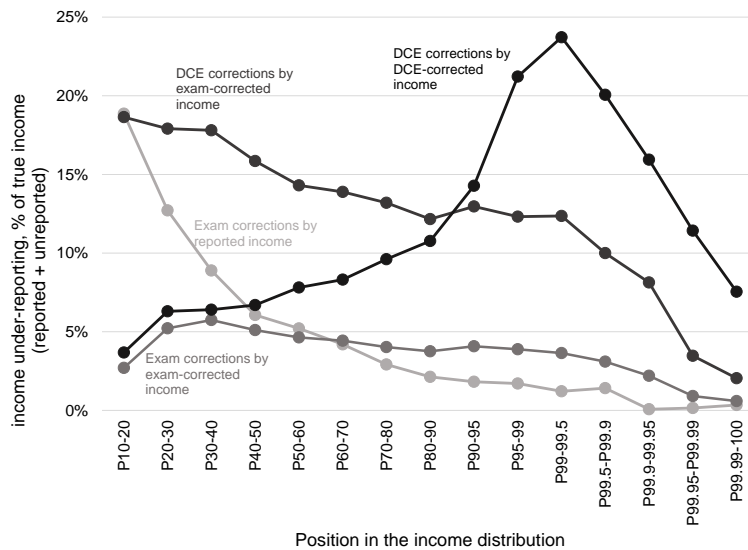
(a) Estimated Income-Under-Reporting



(b) Comparison with Johns and Slemrod (2010)



(c) Illustration of Re-Ranking



Note: This figure illustrates our estimates from applying DCE2001 to NRP random audit data from tax years 2006–2013. See the text and [Johns and Slemrod \(2010\)](#) for a description of DCE2001. Panel (a) presents the estimates of the main statistics of interest to us, which is the share of true income under-reported by rank in the estimated true income distribution. Panel (b) combines some bins at the top of the distribution to make the estimates comparable to earlier work on 2001 NRP data by [Johns and Slemrod \(2010\)](#), showing that with this method we obtain a similar profile of evasion to prior work. Panel (c) unpacks the re-ranking entailed in the allocation method. We start by illustrating the profile of evasion using only exam corrections, i.e. including no undetected under-reporting, and ranking individuals by their position in the reported income distribution, in “exam corrections by reported income.” We then re-rank individuals by their exam-corrected income in “Exam-corrections by exam-corrected income.” Third, we allocate undetected evasion following the multiplier method of DCE2001, but continue to rank by exam-corrected income, in “DCE corrections by exam-corrected income.” Fourth, we re-rank by DCE-adjusted income, i.e. exam-detected income plus the amount of undetected under-reporting allocated to the individual by the method, in “DCE corrections by DCE-corrected income,” to arrive at same estimates as Panel (a).

structural model. More specifically, our understanding is that the dollar amount is based on the conditional expectation of the logarithm of corrected income and an anti-log transform – using an *unconditional* variance for log income, i.e.

$$\hat{y}_{\text{cor}} \equiv \exp \left\{ E[\log(y_{\text{cor}})|X, \Delta y > 0] + \frac{\{Var[\log(y_{\text{cor}})]\}}{2} \right\},$$

where  $E[\log(y_{\text{cor}})|X, \Delta y > 0]$  and  $Var[\log(y_{\text{cor}})]$  are estimates from the structural model.

The simulation then takes random draws to specify whether each taxpayer under-reported each type of income given their estimated conditional probability of under-reporting. As described in IRS (2019), the simulations were run 10 times; following the same approach, we take an average across all 10 simulations in estimating aggregate statistics. We note that  $\hat{y}_{\text{cor}}(X)$  is non-stochastic conditional on the observables on the originally reported return; the simulation does not *not* take a random draw from the estimated conditional distribution of corrected incomes. To obtain an estimated profile of under-reporting, we rank individuals in the simulated data by their rank in the true income distribution, and then estimate distributional statistics.

While the direction of bias in DCE2001 was unclear, there are good ex ante reasons to expect that the simulated DCE2019 profile of under-reporting under-estimates the concentration of undetected under-reporting. The new method corrects the unrealistic feature of the DCE2001 allocation alluded to above by incorporating an extensive margin of under-reporting at the line item level, but it likely over-corrects the problem by loading essentially *all* variation in both undetected *and detected* under-reporting (conditional on observed return characteristics) on the extensive margin of under-reporting in the simulation.<sup>15</sup> Each individual is only allocated a specific dollar amount of under-reporting when under-reporting occurs, which limits the heterogeneity in under-reporting between individuals simulated to be under-reporters. This is likely to lead to under-estimation at the top of the distribution, where the re-ranking attributable to intensive margin variation in under-reporting is mechanically most important – individuals with especially large amounts of under-reporting are more likely to wind up at the top of the true distribution after re-ranking. Additionally, the use of a conditional variance could further bias downwards the likelihood of simulating large corrections for top-income taxpayers.<sup>16</sup>

We report the results from the simulations of the DCE2019 model in Figure 5. The profile of under-reporting from DCE2019 is very different from the DCE2001 estimates, with more under-reporting allocated to the bottom half of the income distribution and less allocated to the top. With this specification, aggre-

<sup>15</sup>We clarify that by “extensive margin” here and elsewhere, we mean the extensive margin of under-reporting with respect to a specific line item of the tax return, such as Schedule C income. If we defined margins in terms of overall income (market income or AGI), there will be some intensive margin variation in under-reporting in the DCE2019 simulation generated by the fact that some taxpayers under-report more than one type of income.

<sup>16</sup>The direction of this bias is ex ante ambiguous and depends on the direction of any heteroskedasticity in log income conditional on observables. We conjecture that this conditional variance is plausibly larger for high-income taxpayers, for example because several well-known forms of evasion (e.g. claiming large erroneous losses) lead to extremely large corrections for a number of these taxpayers.

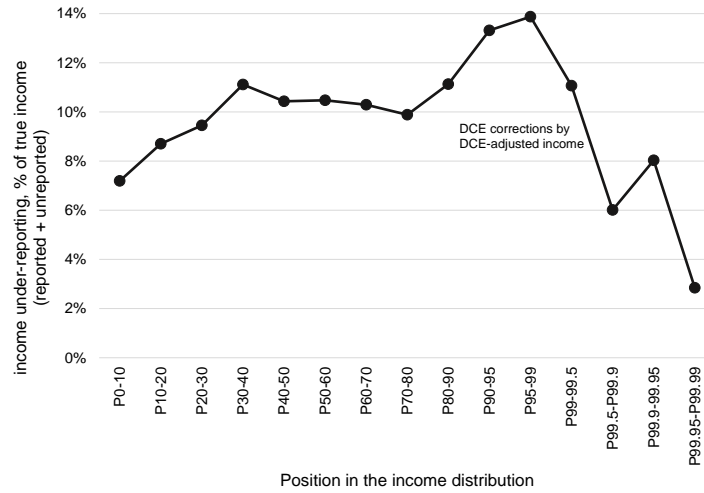
gate under-reporting is spread out across many more individuals, and this flattens the estimated profile of under-reporting compared to DCE2001. This is unsurprising given the discussion above. The pattern of re-rankings is also quite different, which is informative about what this method effectively assumes about undetected under-reporting.

As we have already seen in Figure 1 above, this method implicitly assumes that undetected under-reporting is distributed *like exam-detected under-reporting*. One can understand why by understanding the underlying methodology, and relatedly, by examining how re-ranking shapes the estimates. Methodologically, the simulation behind DCE2019 builds on the fact that a lot of individuals in the bottom half of the reported income distribution are found to be under-reporting by examiners. Because it loads all variation in under-reporting on the extensive margin, the simulation supposes that a lot of individuals in the bottom half, many more than the exam-corrected estimates, are under-reporters and allocates them a typical amount of under-reporting for someone with a tax return like theirs. Re-ranking pushes the individuals simulated to under-report upwards in the distribution compared to their reported income, but, partly because of the limited heterogeneity in intensive margin amounts of under-reporting, it seldom re-ranks them to the top. We end up with a lot more under-reporting in p50-p90 of the estimated true income distribution, compared to DCE2001, and less under-reporting in the top 5%. Figure 5b illustrates that the re-ranking entailed by DCE2019 matches this intuition. We observe a very similar character of re-ranking when transitioning from 1) exam-detected under-reporting by reported income to the same by exam-corrected income, and 2) DCE corrections by reported income to the same by DCE-adjusted income. In both cases, substantial re-ranking initially appears at the bottom of the reported income distribution, and re-ranking pushes this under-reporting to the top half of the distribution, with little under-reporting re-ranking all the way into the top 1%.

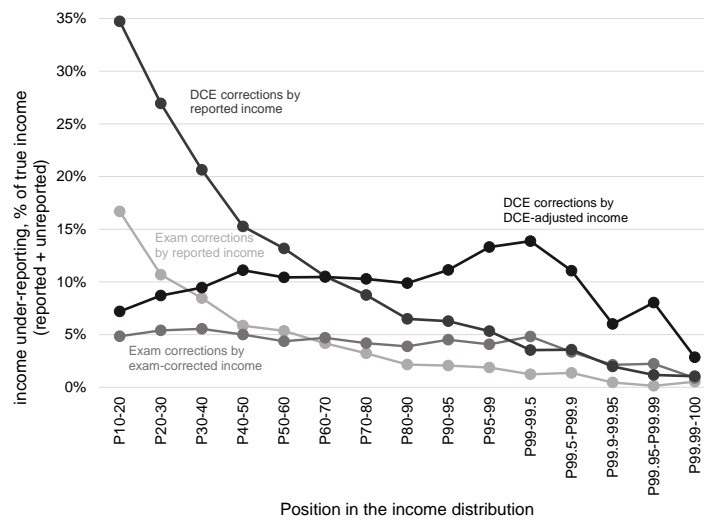
Based on all this, we argue that even if it identifies total undetected under-reporting of income accurately, the DCE2019 allocation of undetected under-reporting likely under-estimates under-reporting at the top of the true income distribution. A simulation aimed at comprehensively modelling the distribution of undetected under-reporting would suppose that 1) there is intensive margin variation in under-reporting and not only extensive margin variation, and 2) sometimes exam corrections do not detect all misreporting on a given return, i.e. there could be detected and undetected under-reporting present on the same return. Incorporating either of these features in the model would cause dollars of under-reporting to re-rank upwards in the income distribution. If we imagine correcting the first issue, individuals simulated to have idiosyncratically larger amounts of mis-reporting would be ranked higher up in the distribution by true income and individuals with smaller amounts would be ranked lower; both of these act to shift under-reporting upward in the distribution. Similarly, accounting for the second issue would cause exam-detected under-reporting

**FIGURE 5: DISTRIBUTIONAL ESTIMATES FOR INCOME UNDER-REPORTING AS A % OF TRUE INCOME USING DCE2019, TY2008–2013**

**(a) Estimated Income Under-Reporting**



**(b) Illustration of Re-Ranking**



Note: This figure reports the results of our application of DCE2019 methods to NRP random audit data from tax years 2008-2013. Panel (a) presents the main estimates. Panel (b) illustrates the influence of re-ranking on these estimates. We illustrate re-ranking in a different fashion from panel (c) of Figure 5 because unlike DCE2001, undetected under-reporting is implicitly allocated at the micro level along with exam-detected under-reporting during the DCE2019 simulations. As such, we report exam corrections by reported and exam-corrected income, as in 4c but for a slightly different set of years, along with DCE corrections by reported and DCE-adjusted income. We argue in the main text that this illustrates a key underlying assumption of DCE2019: that undetected under-reporting is distributed like exam-detected under-reporting.

to be re-ranked upwards in the income distribution when combined with undetected under-reporting.

It is difficult to shed light on this question empirically, but we document one empirical fact suggesting that DCE2019 under-estimates under-reporting at the top of the distribution significantly. Specifically, we compare non-compliance allocated to individuals with high reported incomes by DCE2019 or DCE2001 to non-compliance detected in operational audits of these same high-reported-income individuals. To do this, we compare the assessed tax in operational audits to the unpaid tax assigned to these individuals from the micro DCE methods.<sup>17</sup> Figure 6 shows that operational audits of approximately 10% of individuals in the top 0.01% of the reported income distribution uncovers more unpaid tax than DCE2001 or DCE2019 suggest should exist for all individuals in the top 0.01%.<sup>18</sup> Our estimate from operational audits reflects an extremely conservative bounding exercise: all individuals in the full population receive weight one, we assign audited individuals the amount of tax assessed in their audit, and we assign unaudited individuals *zero unpaid tax*.<sup>19</sup> In other words, we suppose that operational audits detect 100% of evasion: those who are audited have no undetected under-reporting after audit<sup>20</sup> and those who were not audited had zero under-reporting. The operational audit estimate should therefore be an extremely conservative lower bound for the true amount of non-compliance in this sub-population.

The lower bound turns out to be informative in the top 0.01% because of the relatively high audit rate in this sub-population. In fact, we observe that this very conservative lower bound exceeds the non-compliance implied by DCE2019, and approximately equals the non-compliance implied by DCE2001. In other words, operational audits of about 10% of top 0.01% taxpayers uncover more non-compliance than DCE2019 estimates imply for the entire top 0.01% population. This strongly suggests there must be some undetected non-compliance at the very top of the distribution, which is not reflected in the DCE2019 or DCE2001 distributional estimates. Insofar as such non-compliance would be detected in the counterfactual the DCE methods attempt to construct, where all audits are conducted by the most thorough auditors, this suggests that the distribution of non-compliance should be more concentrated at the top of the distribution than DCE2019 implies, which would be consistent with the conceptual limitations of DCE2019 that we discussed above. Further the comparison is much starker with the DCE2019 method than with DCE2001 – making a compelling case that DCE2019 under-estimates the top of the income distribution more than the

---

<sup>17</sup>We do not have comparable data on under-reported income in the operational audit data, so we do this comparison on the basis of under-paid tax. We note that both the DCE estimates and operational audit estimates are based on initial assessments by auditors, so they are comparable in this sense. Likewise we do not have corrected incomes for unaudited individuals in the operational audit data, so we rank individuals by reported income for this comparison. Neither of these issues undermines our interpretation of the results.

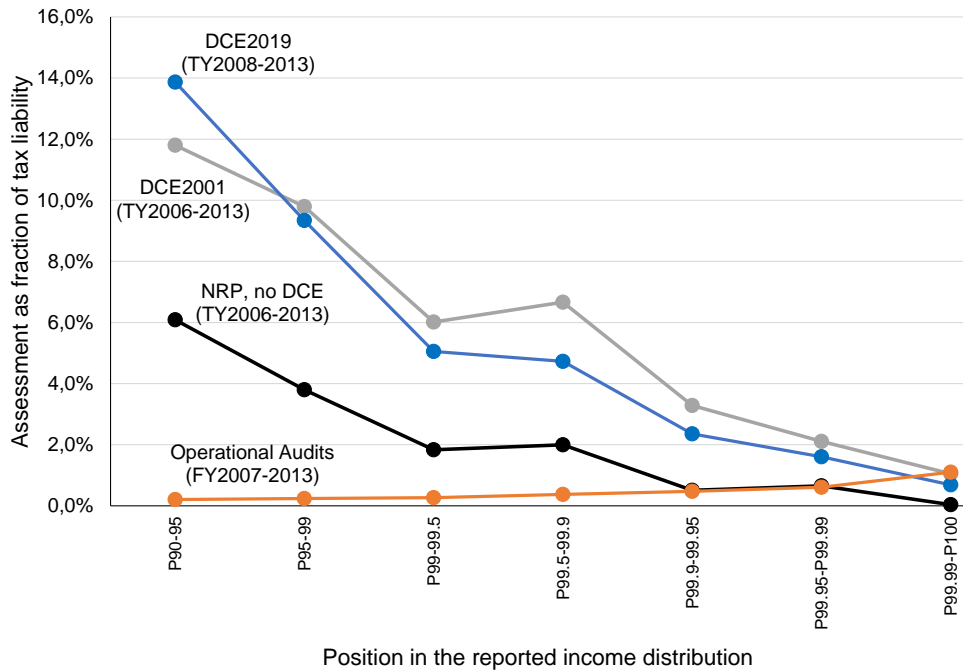
<sup>18</sup>A similar figure appears in GLRRZ; we update the figure to include DCE2019 results here, and we average the operational audit statistics over a number of years. The Appendix of GLRRZ contains some supplementary results that reinforce our claim here.

<sup>19</sup>The important point is that we do not weight observations of operational audits to try and make them representative of the full population; doing so is obviously complicated by the selection of operational audits.

<sup>20</sup>This is highly conservative because, given the overall results from DCE, undetected evasion from individual random audits is often substantial, and the same is plausibly true for operational audits.

DCE2001 method, and that the DCE2001 method despite its conceptual limitations is also likely somewhat downward biased at the top.

**FIGURE 6: COMPARISON OF DCE2001 AND DCE2019 RESULTS WITH OPERATIONAL AUDIT RESULTS AT THE TOP OF THE REPORTED INCOME DISTRIBUTION**



This figure compares the rate of under-payment of taxes in DCE2019, DCE2001, NRP random audits without undetected under-reporting, and operational audit data, ranking taxpayers in the reported income distribution. The y-axis is the ratio of additional taxes assessed by the auditors, i.e. under-paid tax, to total true tax, i.e. under-paid tax plus reported tax from the originally filed Form 1040. For the result based on operational audit data we construct a conservative lower bound, presuming that all individuals not subject to audit in a given year had zero under-payment of taxes in that year. We average the rate of under-payment across tax years 2007 to 2013. Our main interest is in the comparison of these data for the top 0.01% of the distribution, where operational audit rates are high enough that the lower bound is informative. Given that only about 10% of individuals in the top 0.01% are audited in each year and that operational audits may not always uncover all under-payment of taxes, the true amount of under-payment in the top 0.01% of the population should be much higher than the operational audit number appearing in the figure. We observe that DCE2019 and DCE2001 estimates of this quantity, total under-reporting in the top 0.01%, are below or similar to the lower bound we obtain from the operational audit data, suggesting that both methods significantly underestimate under-payment in this group. A similar figure, using operational audit data from 2010 only and without DCE2019, appeared in the first working paper version of GLRRZ.

In summary, evaluating the assumptions imposed by micro methods on the distribution of undetected under-reporting is difficult. Features of the model that are unimportant for the identification of an aggregate amount of undetected under-reporting, such as the extent of within- versus between-individual variation in under-reporting, are essential for distributional estimates. Unfortunately, these distributional questions are not well identified empirically from existing methods. We can imagine several alternative specifica-

tions, and many tweaks to improve the existing specifications (at the risk of potentially also making them intractable or infeasible). Reck et al. (2021) explored some of the feasible alternatives to DCE2001 in illustrating why the direction of the bias in DCE2001 estimates is ambiguous. We ultimately contend that, because the existing methods were not explicitly designed to identify undetected under-reporting throughout the income distribution, existing methods and minor tweaks to them provide limited insights on the distributional questions examined here. What is needed is a return to first principles to design new methods that explicitly aim to identify undetected under-reporting throughout the distribution, which represents a challenging technical problem. We next turn to top-down macro allocation methods, one advantage of which is that they permit us to abstract away from specific features of micro models, and rather illustrate how different assumptions about the concentration of undetected under-reporting directly influence our estimates of interest. We emphasize that these macro allocations should be seen as complementary to the existing micro approaches, as they help us understand the implicit assumptions about the relationship between detection probabilities and true income made by micro methods.

### 4.3 Macro Allocations

We next turn to macro allocation methods, with which we make various explicit and direct assumptions about the location of undetected under-reporting through the true income distribution. A key advantage of this approach is therefore that we can directly focus our attention and assumptions on the quantities that matter for the final allocation of undetected under-reporting. It remains difficult to assess whether these assumptions are realistic, but the assumptions map straightforwardly to distributional estimates, and they relate straightforwardly to theoretical concepts. Relative to a micro allocation, one drawback to the macro approach is that we cannot directly examine re-ranking at the micro level. The extent of re-ranking that occurs when we incorporate undetected under-reporting into the estimates becomes implicit, as the macro allocation assumptions only specify the distribution of undetected under-reporting by true income and not by reported or exam-corrected income.

The basic approach of the macro allocation methods we develop here is to directly specify for each type of income the share of undetected under-reporting attributable to each part of the true income distribution. Although we cannot directly observe where undetected under-reporting should fall, we can use empirical data on types of income that we do observe to guide the specification of these shares. Throughout, we allocate detected under-reporting according to its location in the exam-corrected income distribution using the information in Figure 1 of GLRRZ, which is also plotted on the figures below.<sup>21</sup>

---

<sup>21</sup>Allocating detected under-reporting in this fashion and then adding undetected under-reporting is a conservative choice with respect to the concentration of under-reporting at the top. As discussed above, some individuals with exam-detected under-reporting



The key methodological question here is which empirical income data to use to specify an allocation of undetected under-reporting; this decision directly shapes the resulting profile of under-reporting. One intuitive approach is to assume that undetected under-reporting of each type of income is distributed like reported income for that type of income. That is, if  $X\%$  of reported interest income belongs to the top 1% ranked by reported income, we would assume that  $X\%$  of undetected interest income belongs to the top 1% by true income. We used the same approach in GLRRZ to allocate entity-level under-reporting in partnerships and S corporations.<sup>22</sup> We call this allocation the *reported income share* allocation.

We plot the results from the reported income share allocation in Figure 7, along with some alternatives we describe shortly. To understand what this specification implies about the relationship between detection and true income rank, we contrast the results with the reported income share allocation in Figure 7 with the estimates we would obtain from exam corrections alone, i.e. without including undetected under-reporting. The reported income share allocation implies that the rate of under-reporting is slightly higher in the top 1% than in the bottom 99% of the distribution. This occurs because significant under-reporting is estimated for several types of income that are highly concentrated at the top of the distribution of reported incomes, as discussed above (see Figure 3). We observe from the exam-corrections-only profile in Figure 7 that that exam-detected under-reporting is more concentrated at the bottom of the income distribution than reported income. If detected under-reporting is more concentrated at the bottom relative to reported income, undetected under-reporting must be more concentrated at the top according to these estimates. This is directionally consistent with the various arguments we discussed above suggesting that accounting for undetected under-reporting should increase the concentration of under-reporting.

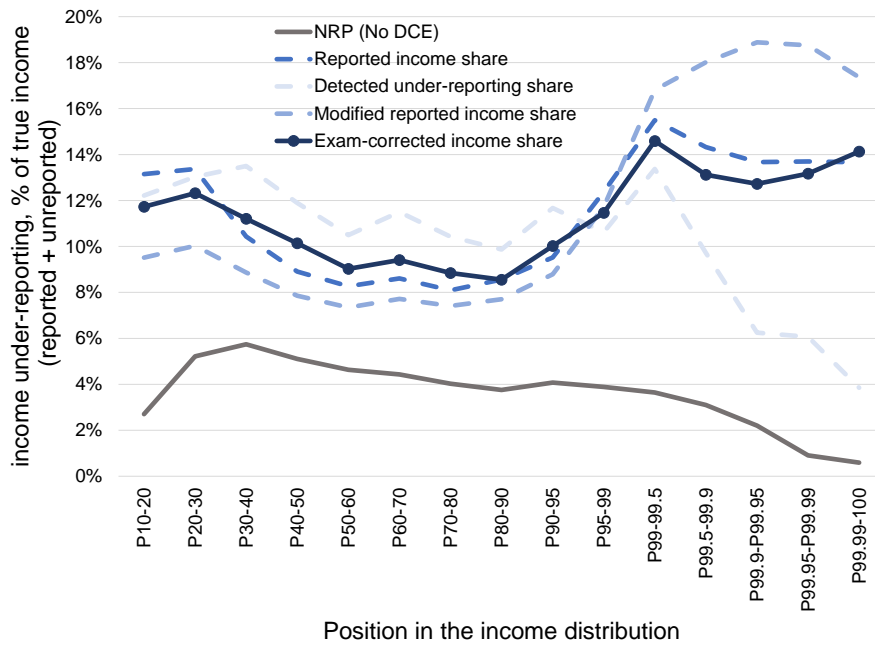
We also report the results of other specifications, which reflect alternative views about the nature of undetected under-reporting. One alternative is to assume that the location of undetected under-reporting resembles that of detected under-reporting, even after-ranking by true income. To construct a specification based on this premise, we can specify the macro allocation shares based on the location of exam-detected under-reporting in the exam-corrected income distribution. That is, if  $X\%$  of exam-detected under-reporting of interest income is allocated to the top 1% by exam-corrected income, then we allocate  $X\%$  of undetected

---

also have undetected under-reporting. Using exam-corrected income to allocate detected under-reporting puts the detected under-reporting by these individuals lower in the income distribution than we would if we could observe undetected under-reporting at the micro level.

<sup>22</sup>As with the partnership allocation, one complication here concerns line items that can be negative. For such line items, we calculate reported income shares based on net amounts within each income bin. When the total for a given line item/type of income is negative in a bin, which only ever occurs in the bottom 10% bin, we re-code this total as zero and then re-calculate the shares, i.e. we allocate zero undetected under-reporting of the given type of income to the bin of the income distribution where total income of the given type is negative. As we explained in Reck et al. (2021), this does not imply we assume individuals with reported incomes in the bottom 10% never under-report these types of income (which would obviously be unrealistic). We do assume, rather, that re-ranking pulls under-reporters in the bottom 10% by reported income out of the bottom 10% when we rank by true income. Based on the mechanics of re-ranking and the extent of re-ranking we typically observe after exam corrections in the NRP data, we believe this assumption to be reasonably appropriate. For example, in NRP data on sole proprietorships (before DCE), we estimate that 14% of all sole proprietorship under-reporting is attributable to taxpayers with reported losses, but just 1.2% of under-reporting is attributable to those for whom exam-corrected income is a loss.

**FIGURE 7: DISTRIBUTIONAL ESTIMATES FOR INCOME UNDER-REPORTING AS A % OF TRUE INCOME USING MACRO ALLOCATION METHODS, TY2008–2013**



Note: This figure plots the results from several different macro allocation methods, to illustrate how varying assumptions about the distribution of undetected under-reporting shape estimates of the overall rate of under-reporting of income. See the main text for descriptions of each of the methods. We contrast each of these with the profile of under-reporting estimated based on exam corrections only, without including any undetected under-reporting.

interest income to the top 1% by true income. We label this the resulting allocation the *detected under-reporting share* allocation and report the results in Figure 7. Exam detected under-reported income in the random audit data is less concentrated at the top than reported income, so this allocation is by nature more conservative than the previous allocation with respect to the concentration of under-reporting at the top and the extent of implicit re-ranking from exam-correction-only estimates.

In fact, we expect this allocation to be too conservative for multiple reasons. First, if under-reporting is higher than exam-corrections-only estimates due to undetected under-reporting, we should expect at least some dollars of under-reporting to re-rank upwards in the distribution when we rank by true income. This allocation implicitly prevents such re-ranking. Additionally, we presented a number of arguments above suggesting that undetected under-reporting should be more concentrated at the top of the distribution than detected under-reporting (see also Phillips, 2014). We therefore regard the detected under-reporting share allocation as an informative lower bound scenario.

There are plausible reasons the reported income share allocation could also under-state the concentration of undetected evasion. As shown by Figure 1, the bulk of the aggregate undetected under-reporting involves business income (sole proprietorships, partnerships, and S-corporations). This makes intuitive sense: while wages, pensions, and most financial capital incomes are subject to extensive third-party reporting, a significant fraction of business income is not. There is, however, an imbalance in the adjustment for undetected under-reporting for the various forms of business income: the sole proprietorship adjustment (\$238 billion on average over 2008-2013, in \$2012) is much larger than the pass-through business income (partnerships, S-corporations) adjustment (\$59 billion). While it is certainly possible that sole proprietorships have more undetected non-compliance than S-corporations and partnerships, another interpretation of this imbalance is that it reflects the asymmetric examination of business income in the NRP. As discussed in GLRRZ, while sole proprietors are comprehensively examined, partnerships and S-corporations (some of which are large multi-state or even multinational businesses) sometimes are not, due to the audit procedures and resource constraints of the random audit program. When certain types of businesses are rarely comprehensively examined, even the most thorough auditors may find little non-compliance where there actually is such non-compliance, leading to biased estimates of undetected under-reporting.

For aggregate estimates of underreported income, whether undetected non-compliance is located on Schedule C (sole proprietorships) or schedule E (partnerships, S-corporations) of the individual income tax return does not matter. However, it does matter for the distribution of noncompliance, because the ownership of pass-through businesses is much more concentrated at the top of the income distribution than the ownership of sole proprietorships. This issue is particularly relevant in light of recent findings that a significant fraction of pass-through business income in the top 1% and top 0.1% of the income distribution

derives from cash-intensive businesses with little third-party information reporting.<sup>23</sup>

To illustrate the relevance of this point, we report the results of a specification in which we treat all forms of business income – sole proprietor income, partnership business income, and S corporation business income – as a single type of income. In other words, if the top 1% by reported income reports  $X\%$  of combined business income, we allocate  $X\%$  of undetected under-reporting of combined business incomes to the top 1% by true income. The underlying assumption is that the aggregate estimate for undetected business income provides a reliable estimate of the amount of undetected under-reported business income of all types (due, e.g., to evasion in cash-intensive industries). Other forms of income are allocated by reported income shares as in the original reported income share allocation. We label this the *modified reported income share* allocation and report the results for this specification in Figure 7. Adopting this specification increases the concentration of undetected under-reporting somewhat because it implicitly reallocates some non-compliance originally allocated to sole proprietorships to more concentrated forms of business income. One could also argue, along similar lines, that undetected under-reporting of financial capital income should also be more concentrated than reported financial capital income because the limitations of third-party information reporting on financial capital income are more pronounced for high-wealth individuals with sophisticated financial arrangements. An additional adjustment along these lines would further increase the estimated concentration of under-reporting.

A final alternative is to allocate undetected under-reporting using exam-corrected income shares. Rather than assuming undetected under-reporting of a given type is distributed like reported income, here we would assume that it is distributed like exam-corrected income, i.e. reported income plus the adjustments made by examiners. For example, if  $X\%$  of total exam-corrected interest income belongs to the top 1% by exam-corrected income, then we would allocate  $X\%$  of undetected under-reporting of interest income to the top 1% by true income. We call this the *exam-corrected income share* allocation and report the results in Figure 7. We observe that the results are very similar to the reported income share allocation. Exam-corrected income distribution is slightly less concentrated than the reported income distribution, so the exam-corrected income share specification leads to slightly lower concentration than the reported income share, but the difference is small.

We note that one property of the exam-corrected income share specification is that it is neutral with respect to the distribution of income of various types. For all of these specifications, we start from the exam-corrected income distribution and then allocate undetected under-reporting. Mechanically, distributing undetected under-reporting like exam-corrected income keeps the overall share of income of each type

---

<sup>23</sup>Smith et al. (2019, Table 3) show that automobile dealers are the second largest source of S-corporation income in the top 0.1% of the income distribution in 2014; specialty trade contractors (pouring concrete, site preparation, plumbing, painting, electrical work...), various professional and technical services, and offices of dentists and physicians also feature prominently.

accruing to each part of the distribution constant. Building on our previous example, if  $X\%$  of total exam-corrected interest income belongs to the top 1% by exam-corrected income, and then we allocate  $X\%$  of undetected interest income to the top 1%, then in the resulting income distribution,  $X\%$  of total interest income belongs to the top 1%. This logic implies that the exam-corrected income share allocation is nearly neutral with respect to the overall concentration of income – meaning e.g. that allocating undetected under-reporting in this way has little impact on the top 1% income share, corrected for under-reporting. The only reason the exam-corrected income share allocation is not perfectly neutral in this sense is that the composition of income changes when we add undetected under-reporting, as discussed above.

#### 4.4 Comparison of Key Aggregates for Various Methods

We next compare all the various methods we considered above, considering what they imply about the profile of evasion and distribution of income. Doing so allows us to understand the relationship between micro and macro methods, and how different views about the distribution of undetected under-reporting shape statistics on inequality.

We now return to Figure 1, in which we compared all of the specifications for estimated income under-reporting by rank in the income distribution that we have considered so far. We observe that in most specifications (all except DCE2001), the rate of under-reporting is flat at around 7-12% of true income for most of the distribution, and then rises as we move into the top 10% of the distribution. The alternative specifications diverge within the top 1% of the distribution, with the detected under-reporting share allocation and DCE2019 implying a sharp drop-off in under-reporting, while the reported and exam-detected income share specifications suggest a tiny decrease or levelling off of under-reporting within the top 1%, and the modified reported income share specification increases under-reporting further within the top 1%. In contrast to all of these, the DCE2001 specification increases throughout the distribution and especially sharply in the P90-99 range, and then drops off sharply within the top 1%.

One important lesson we learn from this comparison is that the detected under-reporting share macro method and the DCE2019 micro method yield very similar results, which confirms our assessment of the DCE2019 method above. Both of these methods assume, implicitly or explicitly, that most of the relevant variation in detection is on the extensive margin, so that undetected under-reporting is distributed like detected under-reporting *even after re-ranking*. We discussed above why both of these methods are likely too conservative, for related reasons. For example, the issue with the absence of intensive margin variation in under-reporting in DCE2019 leads to limited re-ranking, while the detected under-reporting share allocation also implies too little re-ranking because some undetected under-reporting should be allocated to

individuals with detected under-reporting.

Another important lesson from Figure 1 is that the sharp increase and decrease in the rate of under-reporting in DCE2001 is difficult to rationalize from macro assumptions about the nature of undetected non-compliance. Given what we also saw about how re-ranking influences DCE2001 estimates, this suggests that the very sharp increase and decrease in the rate of under-reporting around the 99th percentile of income in DC2001 are not credible. However, if we imagine averaging the DCE2001 profile of under-reporting across bins in the bottom 99%, and separately across bins within the top 1% of the distribution, we would obtain a profile similar to the reported income or exam-corrected income share allocations. In this sense, the reported income share allocation resembles a “flattened out” version of DCE2001. Keeping this fact in mind is useful for interpreting the next set of results, on what these alternative specifications imply about the concentration of under-reported and total income in the top 1% of the distribution.

Next we consider how the distribution of undetected evasion influences overall inequality statistics, like the fraction of under-reporting attributable to the top 1% of the distribution and the top 1% income share corrected for under-reporting. Table 3a reports the share of unreported income attributable to different parts of the estimated true income distribution, according to all the different specifications we have considered, averaging over tax years 2008–2013. To facilitate comparison, we begin with exam-corrected data without accounting for any undetected under-reporting. The first column of Table 3a reports the share of detected under-reporting accruing to different parts of the income distribution, ranking individuals by exam-corrected income. We estimate that 12% of detected under-reporting over tax years 2008–2013 belongs to the top 1% of the exam-corrected income distribution. All methods for allocating undetected under-reporting increase the top 1% share of under-reporting, but to a greatly varying extent. The DCE2001 specification suggests 24.4% of under-reporting belongs to the top 1% of the true income distribution; the reported income share macro allocation estimate of the same quantity is similar, at 26.3%. Both of these are also similar to the results of [Johns and Slemrod \(2010\)](#), who estimated that 27% of under-reporting belongs to the top 1% by true income using DCE2001 methods on tax year 2001 data. As in Figure 1, DCE2019 and the detected under-reporting share macro allocation imply 15.3% and 14.6% of under-reporting belongs to the top 1%, respectively. These estimates are similar to one another and, unsurprisingly, substantially less concentrated than the DCE2001 or reported income share macro estimate. The exam-corrected income share allocation resembles the reported income share allocation, with 24.9% of under-reporting in the top 1%. At the top of this range of estimates of concentration of under-reporting, the modified reported income share allocation suggests that 34.3% of under-reported income belongs to the top 1%.

We next consider the estimation of corrected top 1% fiscal income shares, accounting for under-reporting of incomes on tax returns. Table 4a reports estimates of the share of income accruing to different parts

of the income distribution for tax years 2008–2013. We begin with the concentration of reported income, ranking by reported incomes, i.e. not accounting for any under-reporting. Importantly for the comparison of subsequent results; The top 1% reported income share is 19.1% on average over 2008–2013.

The second column of Table 4a incorporates exam-detected under-reporting into total income and ranks individuals by exam-corrected income, i.e. we estimate the distribution of exam-corrected income without including undetected under-reporting. As exam-corrected under-reporting is more concentrated at the bottom 99% of the distribution than reported incomes, incorporating exam-detected under-reporting alone decreases the top 1% income share by 0.4 percentage points.

Adding undetected under-reporting in the next several columns, we observe that the allocation of undetected under-reporting exerts significant influence on the corrected top 1% income share. As in the previous table, the DCE2019 micro method and the detected under-reporting share macro allocation yield similar estimates of the corrected top 1% income share, at 18.1% or 18.4% respectively. Both of these specifications further decrease the concentration of income, because they assume that, like exam-detected under-reporting, undetected under-reporting is less concentrated in the top 1% than reported or exam-corrected incomes. The DCE2001 estimates imply a top 1% share of 19.1%, implying that allocating (detected and undetected) under-reported income has a net zero effect on the top 1% income share, as in Johns and Slemrod (2010). The macro reported income share allocation yields modestly larger estimates, with a corrected top 1% income share of 19.7%. Unsurprisingly, the modified reported income share allocation leads to the highest concentration of income, implying a corrected top 1% income share of 20.5%.

The exam-corrected income share macro allocation estimate of the top 1% income share is 19.5%. This is similar to but slightly smaller than the macro reported income share allocation estimate (19.7%), and slightly higher than the estimate based on reported incomes without under-reporting (19.1%). It is especially informative to compare the exam-corrected income share macro allocation estimate to the estimate based on exam corrections only in the second column of Table 4a (18.7%). As discussed above, the macro exam-corrected income share allocation is distributionally neutral by type of income, relative to the exam-corrections-only estimates in second column of Table 4a. As a result difference between these two estimates, implying an increase in the top 1% share of 0.8 pp, is entirely driven by the change in the composition of income when we incorporate undetected under-reporting. As discussed above, estimated undetected under-reporting largely consists of types of income that are concentrated at the top of the distribution, so the change in the composition of income between these estimates increases the estimated top 1% income share somewhat.

Tables 4b and 4b report the same information as Tables 3a and 4a, but averaging over the slightly wider set of years in GLRRZ, tax years 2006–2013 rather than 2008–2013. To do this, we assume that undetected

under-reporting of each type of income, expressed as a share of total true income (i.e. the exact data plotted in Figure 2) was the same in 2006–2013 as in 2008–2013. We do not include DCE2019 here as we do not have DCE2019 micro simulations covering this set of years, but we infer from the above that the results would be similar to the detected under-reporting share macro allocation estimates. The main difference between these and the results discussed above is that the years of the financial crisis are more heavily weighted in the 2008–2013 estimates than in the 2006–2013 estimates. The top 1% income share fell during the crisis, so income is slightly less concentrated using 2008–2013 than using 2006–2013. For example, the top 1% income reported income share is 1.2 percentage points lower in 2008–2013 than in 2006–2013. Virtually all the differences between different methods for allocating undetected under-reporting are the same in Tables 4b and 4b as in the set of results discussed above.

In summary, allocating *detected* under-reporting based on random audit data decreases the top 1% income share by about 0.4 pp relative to reported incomes only. Further allocating *undetected* under-reporting modifies the top 1% income share by -0.6 to +1.8 pp across all specifications we consider, relative to results with only detected under-reporting. Insofar as we regard total income shares as the main statistics of interest in the study of inequality, the most justifiable allocation of undetected under-reporting in our view is one that does not cause a sizable revision of the top 1% share in either direction. We have no direct evidence on the location of undetected under-reporting in the distribution because it is by nature unobserved. All the feasible allocation methods we have considered here require assumptions, the validity of which is difficult to test empirically. Inequality could be higher or lower due to undetected under-reporting, and we do not have enough information to be certain about this. As such, the most neutral approach is arguably to avoid revising the top 1% income share substantially in either direction when accounting for under-reporting. Three of the specifications for distributing undetected under-reporting we have considered have this property: the reported income share and exam-corrected income share macro allocations, and DCE2001. However, we do not recommend using DCE2001 to distribute undetected evasion because, as discussed above, the profile of under-reporting *within* the top 1%, and within the bottom 99%, seems unrealistic compared to the alternatives we considered.

Further, the increase in the top 1% share from the exam-corrected income share allocation is entirely driven by changes in the composition of income, occurring because estimated undetected under-reporting largely consists of types of income that are highly concentrated at the top of the distribution. In our view, this property-maintaining distributional neutrality for each type of income but accounting for changes in the composition of total income-makes the exam-detected under-reporting share macro allocation the most defensible method among those we have considered. However, we caution that this does not imply that this specification is absolutely definitive, and we hope that future research will help to refine the estimates



and improve our understanding of the distribution of undetected under-reporting.

## 5 Unconfronted Methodological Questions

Our main focus here has been the distribution of undetected under-reporting of income. As such, throughout the above analysis we essentially took the DCE estimates for total under-reporting of each type of income in the full population at face value, setting aside some other issues with DCE methodology. We review these methodological issues here. Some of these issues are also discussed in a recent paper by [Hemel et al. \(2021\)](#).

One concern is that the auditors working NRP cases may be assigned in a selected fashion. The efficient case assignment rule from a tax collections perspective would arguably be to assign the most thorough and experienced auditors to cases suspected of containing the most difficult-to-detect evasion. If the existing assignment rule is relatively efficient in this respect, this could cause bias in DCE estimates. If the most capable auditors tend to be assigned to the cases with the most non-compliance, then DCE estimates would tend to over-estimate total under-reporting.

Second, these estimates are based on initial auditor assessments. Pondering that these initial assessments may be imperfect raises a number of issues. One possibility is that the auditors with the highest measured detection rates are over-aggressive, i.e. they “detect” more than 100 percent of true under-reporting on the tax returns they audit. In this case, DCE totals would over-estimate total true evasion.

This second concern actually scratches the surface of an even deeper conceptual challenge. The underlying issue here is that the auditor’s initial recommended assessment reflects in part a judgment or interpretation of tax rules, while an objective judgment about whether some tax position is legal is rarely if ever observed. To better understand this challenge, suppose for simplicity there are just two auditors, Gabriel and Daniel, and Gabriel systematically recommends larger assessments than Daniel. This could actually happen for two basic reasons. The preferred explanation implicit in current DCE methods is that one auditor systematically reviews more information, which leads to more accurate assessments. From there, the logic of the model is essentially that if taxpayers generally conceal information related to *under-reporting*, reviewing more information will also lead to larger assessments. By this logic Gabriel systematically finds more under-reporting, so his assessments will be more accurate and thorough than Daniel’s. Thus we should try to construct a counterfactual wherein Gabriel performed all the audits. But this logic may not obtain universally, e.g. with more than two auditors the auditors with the very highest rates of assessment may not be the very most thorough and accurate. Another reason that auditors could differ systematically is that they reviewed basically the same information, but they interpret the rules systematically differently. In this case, the DCE method presumes that Gabriel’s interpretations of the rules are correct because he

**TABLE 3A: COMPARING ALLOCATION METHODS: UNDER-REPORTED INCOME AS A % OF TOTAL UNDER-REPORTED INCOME - TY2008-2013**

Bin	Only detected under-rep.	DCE2001	DCE2019	Reported inc. share (macro)	Detected under-rep. share (macro)	Modified rep. inc. share (macro)	Exam-corr. inc. share (macro)
P0-10	-0.7	-0.5	0.9	-0.1	0.7	-0.1	-0.2
P10-20	0.9	0.3	1.0	1.9	1.7	1.3	1.6
P20-30	1.7	1.0	2.1	3.2	3.1	2.3	2.8
P30-40	2.5	1.6	3.3	3.5	4.7	2.9	3.6
P40-50	3.1	2.4	5.2	4.0	5.5	3.5	4.4
P50-60	3.7	3.9	6.4	4.9	6.3	4.3	5.1
P60-70	5.5	5.6	8.4	6.7	9.2	5.9	7.1
P70-80	7.0	8.5	11.0	8.4	11.2	7.7	8.9
P80-90	9.8	14.0	14.9	12.7	14.9	11.3	12.3
P90-95	9.1	13.1	12.1	10.2	12.8	9.4	10.5
P95-99	12.0	25.6	19.3	18.3	15.4	17.2	16.5
P99-99.5	4.1	8.3	5.1	6.0	5.1	6.6	5.6
P99.5-99.9	4.6	10.6	6.2	8.7	5.6	11.4	7.9
P99.9-P99.95	0.9	2.1	1.0	2.4	1.0	3.5	2.2
P99.95-P99.99	1.5	2.5	2.0	3.8	1.6	5.6	3.7
P99.99-100	0.9	1.0	1.0	5.4	1.4	7.2	5.6
Top 1%	12.0	24.4	15.3	26.3	14.6	34.3	24.9

Note: This Table presents estimates of the distribution of under-reporting through the income distribution, on average over tax years 2008-2013. In each column, we report the share of total under-reporting of income attributable to different bins of estimated true income. Different columns reflect different specifications. We begin in the first column with a specification including only exam-detected under-reporting and no undetected under-reporting. We include this first column for comparison purposes; all other columns include detected and undetected under-reporting. The next two columns specify undetected under-reporting according to DCE2001 and DCE2019. We then report estimates using macro allocation methods in the final four columns.

TABLE 4A: COMPARING ALLOCATION METHODS: TRUE INCOME AS A % OF TOTAL TRUE INCOME - TY2008-2013

Bin	Reported incomes	Exam-corr. incomes	DCE2001	DCE2019	Reported inc. share (macro)	Detected under-rep. share (macro)	Modified rep. inc. share (macro)	Exam-corr. inc. share (macro)
P0-10	-1.3	-1.0	-1.9	-0.6	-0.8	-0.8	-0.8	-0.9
P10-20	1.4	1.5	1.0	1.5	1.5	1.5	1.4	1.5
P20-30	2.4	2.5	2.1	2.6	2.6	2.5	2.5	2.5
P30-40	3.5	3.6	3.3	3.7	3.6	3.7	3.5	3.6
P40-50	4.8	4.9	4.7	5.0	4.8	4.9	4.7	4.8
P50-60	6.4	6.5	6.4	6.5	6.3	6.4	6.2	6.3
P60-70	8.5	8.5	8.6	8.6	8.3	8.6	8.2	8.3
P70-80	11.5	11.5	11.6	11.4	11.1	11.4	11.0	11.2
P80-90	16.3	16.2	16.5	16.1	15.9	16.1	15.7	15.8
P90-95	11.8	11.7	12.1	11.6	11.5	11.7	11.4	11.5
P95-99	15.6	15.5	16.5	15.5	15.8	15.5	15.7	15.6
P99-99.5	4.0	4.0	4.3	3.9	4.1	4.0	4.2	4.1
P99.5-99.9	6.3	6.1	6.5	6.0	6.4	6.1	6.7	6.4
P99.9-P99.95	1.8	1.8	1.8	1.7	1.9	1.7	2.0	1.8
P99.95-P99.99	2.9	2.8	2.8	2.7	3.0	2.7	3.2	3.0
P99.99-100	4.1	4.0	3.8	3.8	4.3	3.8	4.4	4.3
Top 1%	19.1	18.7	19.1	18.1	19.7	18.4	20.5	19.5

Note: This Table presents estimates of the distribution of total income, averaging over tax years 2008–2013. In each column, we report the share of total income attributable to different bins of estimated true income. Different columns reflect different specifications. We begin in the first column with reported incomes, ranking individuals by their position in the reported income distribution. We add detected under-reporting only in the next column, estimating the distribution of exam-corrected incomes, ranking individuals by their exam-corrected income. We include these first two columns for comparison purposes; all other columns include detected and undetected under-reporting in total income. Following these two, the in next two columns we specify undetected under-reporting according to DCE2001 and DCE2019 and then estimate total income shares. We report estimates using macro allocation methods in the final four columns.

systematically assesses more. But which auditor is more often correct may in fact be ambiguous. We might argue, as before, that where there are ambiguities in the rules, taxpayers are likely to err on the side of paying less tax, so that, on balance, a more accurate take on the rules would find more under-reporting. But once again this may not be universally true throughout the distribution of auditors.

A more structural interpretation of differential detection along these lines is that the legality of a given position is fundamentally uncertain, and different auditors have different thresholds for recommending an assessment. To think through this, suppose that each potential adjustment indicating under-reporting has an observable probability of being reversed in appeals, and this probability is common knowledge to auditors (reflecting that they reviewed the same information, etc). Suppose further that the auditors recommending the largest and most frequent assessments used a 10% threshold, so that they recommended additional assessments that would hold up on appeals 10% of the time. Meanwhile suppose that on average auditors used a threshold around 50%, so they recommended systematically fewer and smaller adjustments than the 10% auditors. Within this type of model, *the definition of the tax gap is itself ambiguous*: should we define the tax gap as 50% under-reporting? As 10% under-reporting? The answer is not clear, but this reasoning at least delivers a clear answer as to what current DCE methods are doing: positions that are relatively more likely to be reversed on appeals would be included in the tax gap after DCE. More fundamentally, DCE estimates reflect the interpretations of tax law that are least favorable to the taxpayer among all auditors' interpretations. Which interpretation is the right interpretation to use in estimating the tax gap is in fact a judgment call, and DCE reflects a relatively harsh judgment (but one which might nevertheless still be appropriate).

One way forward suggested by [Hemel et al. \(2021\)](#) would be to simply observe whether an assessment is reversed on appeal by the taxpayer to resolve these ambiguities. We agree that additional research on final outcomes of cases would be informative, but deferring to final outcomes in estimating the tax gap is not as perfect a solution as one might expect. Whether a case goes to appeals is endogenous. Not everyone has the time, resources, or motivation to dispute their assessments; protracted disputes are especially costly. Figure 8, drawn from GLRRZ, shows that higher-income individuals are far more likely to contest their assessments than others. This figure depicts just the first stage of the disagreement and appeals process, but at each stage, those with more time and resources to fight an assessment will tend to do so. The resources of the IRS are likewise limited. The vast majority of cases settle without a judge weighing in on the correct interpretation of the law. In the end, the final outcomes of cases therefore reflect a combination of the correction of over-assessments by auditors *and taxpayers' capacity to engage with the tax system in order to dispute their assessments*. Without accounting for the latter, deferring to the "final outcome" of a case would likely bias our estimates of the tax gap downwards, especially at the top of the distribution.

There are a few more technical concerns about how DCE methods are actually implemented as well. One concern is that sampling variation in estimated auditor effects would lead to a distribution of auditor fixed effects that is overly dispersed (Efron and Morris, 1975). Intuitively, each estimated auditor effect contains some noise, so the very top of the auditor effects distribution may be due to some auditors who were, by chance, assigned cases containing large amounts of under-reporting. If this type of sampling variation influences the estimate of the 100% detection counterfactual, it will tend to cause DCE to over-estimate total evasion. There are well-established estimators available to deal with this issue (Efron and Morris, 1975); it is unclear how much of a difference this would make. To mitigate the problem, researchers do restrict the estimation of the DCE model to a subset of auditors taking on a significant number of cases.

Another technical concern is that implementing DCE requires grouping certain types of income together in estimation. The IRS currently estimates separate detection processes for three categories of income, essentially based on the extent of third-party information reporting supporting a given type of income (see table 4 of IRS, 2019). This grouping amounts to a restriction on the underlying model; some restriction along these lines is necessary to obtain precise estimates. Inherently, this practical concern implies that there is more uncertainty about rates of under-reporting of specific types of income with DCE than there is uncertainty about overall rates of under-reporting of income. The delineation of these categories raises some concerns, especially at the top of the distribution, where many taxpayers have control over what is reported on forms we might elsewhere think of as “third party” reports. For example, individuals with a controlling stake in a pass-through business can control what is reported on their Schedule K-1, while for other owners, the Schedule K-1 better conforms with the definition of a third-party report. Such concerns are most prominent near the top of the income distribution, where e.g. controlling owners tend to locate, which could create downward bias in both the overall level and concentration of undetected misreporting.

Another technical concern involves the stratified sampling and weighting procedure. In an effort to increase coverage of those returns likely to contain the most non-compliance (in dollar terms), the random audit program over-samples high-income individuals and individuals with self-employment income. However, the sampling process can condition only on information reported on individuals tax return. As such, researchers can and do over-sample returns with large reported incomes but they cannot directly over-sample returns with large *true* incomes. Moreover, the sampling rates for the least risky sub-groups are relatively low, so that in weighted terms one individual can effectively represent thousands of others in population-weighted statistics. As a result, estimates of the total tax gap or total income under-reporting can be far more highly leveraged than one might naively expect, especially at the top of the distribution when ranking by exam-corrected income. In a small number of cases, examiners find that an individual appearing based on their reported tax return to be relatively low-income and low-risk turns out to be massively

non-compliant. Because very few such individuals are sampled, this introduces more noise than one might expect in estimates of overall non-compliance. For example, an observation with a sampling probability of 0.001 effectively represents 1,000 people, so if one such person is found to have \$1,000,000 of unreported income, this observation could contribute over \$300 million to the estimated total tax gap. One might wish to sample such individuals at a higher rate to reduce the resulting noise, but there is no obvious way to do this.

Finally, in light of all of the uncertainty here, we can understand why some readers may wish to give up on DCE, at least for distributional analysis. We reported results in GLRRZ so that readers who wish to adopt this perspective could do so and still appreciate the main argument of GLRRZ regarding the measurement of sophisticated evasion at the top of the distribution. Nevertheless, we caution readers against giving up on DCE too readily. Despite the difficulty inherent in its estimation, the bulk of the evidence suggests that undetected misreporting is materially present across the entire income distribution. While difficult to perfect, DCE methods give us some way of assessing how large this issue is likely to be, and the estimated models indeed signal that quantifying undetected under-reporting is very important. We therefore hope that our critical discussion here be taken in a constructive spirit, which might inform and inspire future work to refine DCE methods, especially for distributional questions.

## References

- Allingham, M. G. and Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics*, 1(3-4):323–338.
- Alstadsaeter, A., Johannesen, N., and Zucman, G. (2019). Tax Evasion and Inequality. *American Economic Review*, 109(6):2073–2103.
- Auten, G. and Splinter, D. (2021). Comment: Tax evasion at the top of the income distribution: Theory and evidence. Comment, available online at <http://www.davidsplinter.com/AutenSplinter-TaxEvasion.pdf>.
- Bloomquist, K., Emblom, E., Johns, D., and Langetieg, P. (2012). Estimates of the tax year 2006 individual income tax underreporting gap. 2012 irs-tpc research conference, publication 1500, Washington, DC.
- DeBacker, J., Heim, B., Tran, A., and Yuskavage, A. (2020). Tax noncompliance and measures of income inequality. *Tax Notes*, February 17, 2020.

- Efron, B. and Morris, C. (1975). Data analysis using stein's estimator and its generalizations. *Journal of the American Statistical Association*, 70(350):311–319.
- Erard, B. and Feinstein, J. (2011). The individual income reporting gap: What we see and what we don't. 2011 irs-tpc research conference, publication 1500 (rev 4-2012), Washington, DC.
- Feinstein, J. S. (1991). An Econometric Analysis of Income Tax Evasion and Its Detection. *The RAND Journal of Economics*, 22(1):14–35.
- Guyton, J., Langetieg, P., Reck, D., Risch, M., and Zucman, G. (2021). Tax evasion at the top of the income distribution: Theory and evidence. Working paper no. w28542, march 2021., National Bureau of Economic Research.
- Hemel, D. J., Holtzblatt, J., and Rosenthal, S. (2021). The tax gap's many shades of gray. Working paper, SSRN.
- IRS (2007). Tax year 2001 federal tax gap overview. <https://www.irs.gov/pub/irs-soi/01rastg07map.pdf>.
- 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.
- Johns, A. and Slemrod, J. (2010). The Distribution of Income Tax Noncompliance. *National Tax Journal*, 63(3):397–418.
- 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):651–692.
- Phillips, M. D. (2014). Deterrence vs. gamesmanship: Taxpayer response to targeted audits and endogenous penalties. *Journal of Economic Behavior & Organization*, 100(4):81–98.
- Piketty, T. and Saez, E. (2003). Income inequality in the united states, 1913–1998. *The Quarterly journal of economics*, 118(1):1–41.
- Reck, D., Risch, M., and Zucman, G. (2021). Response to a comment by Auten and Splinter on “tax evasion at the top of the income distribution: Theory and evidence”. Mimeo, available at [https://www.danreck.com/s/asresponse\\_full.pdf](https://www.danreck.com/s/asresponse_full.pdf).

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

## **A Additional Results**



**TABLE 4B: COMPARING ALLOCATION METHODS: UNDER-REPORTED INCOME AS A % OF TOTAL UNDER-REPORTED INCOME - TY2006-2013**

Bin	Only detected under-rep.	DCE2001	Reported inc. share (macro)	Detected under-rep. share (macro)	Modified rep. inc. share (macro)	Exam-corr. inc. share (macro)
P0-10	1.4	-0.1	0.9	3.2	0.9	0.8
P10-20	0.7	0.3	1.5	0.8	0.9	0.8
P20-30	2.8	1.0	3.0	2.7	2.1	2.5
P30-40	4.9	1.6	3.5	4.6	2.9	3.9
P40-50	6.2	2.4	4.0	5.9	3.5	4.8
P50-60	7.6	3.7	5.1	7.1	4.5	6.2
P60-70	9.7	5.3	6.6	9.0	5.9	8.4
P70-80	11.9	8.3	8.5	11.3	7.7	11.0
P80-90	15.6	13.2	12.8	15.3	11.4	14.7
P90-95	12.2	12.8	10.0	12.5	9.1	11.9
P95-99	15.7	26.5	18.4	16.1	17.3	18.5
P99-99.5	3.9	7.8	5.7	4.1	6.3	5.7
P99.5-99.9	5.1	10.4	8.7	6.0	11.4	8.3
P99.9-P99.95	1.1	2.4	2.4	1.2	3.6	2.3
P99.95-P99.99	0.7	2.4	3.5	0.7	5.3	3.2
P99.99-100	0.6	2.1	5.4	0.9	7.1	5.4
Top 1%	11.4	25.1	25.7	12.9	33.7	24.8

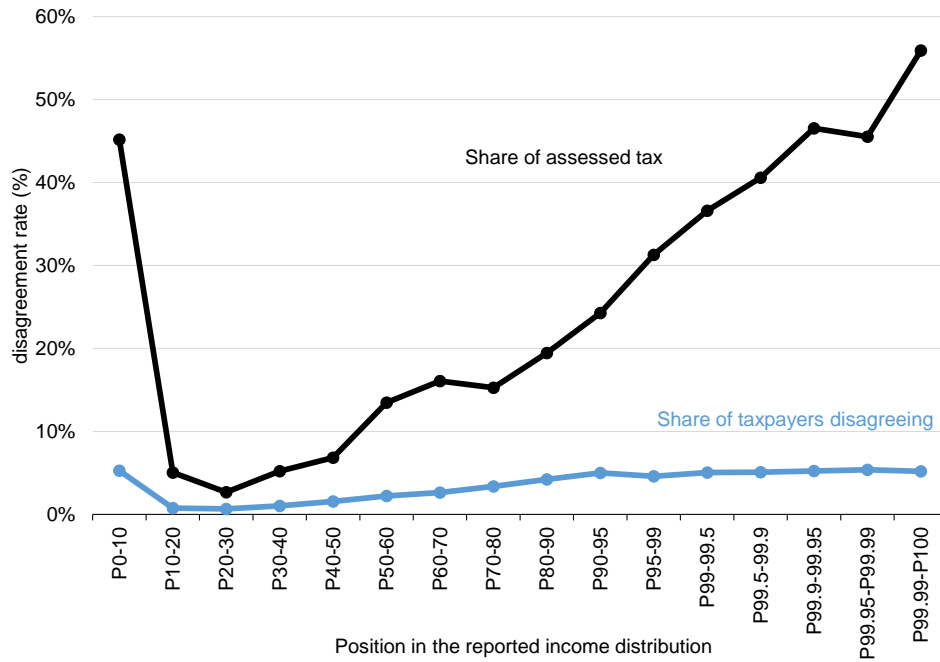
Note: This Table presents estimates of the distribution of under-reporting through the income distribution, on average over tax years 2006–2013. In each column, we report the share of total under-reporting of income attributable to different bins of estimated true income. Different columns reflect different specifications. We begin in the first column with a specification including only exam-detected under-reporting and no undetected under-reporting. We include this first column for comparison purposes; all other columns include detected and undetected under-reporting. The next column specifies undetected under-reporting according to DCE2001. We then report estimates using macro allocation methods in the final four columns.

TABLE 4B: COMPARING ALLOCATION METHODS: TRUE INCOME AS A % OF TOTAL TRUE INCOME - TY2006-2013

Bin	Reported incomes	Exam-corr. incomes	DCE2001	Reported inc. share (macro)	Detected under-rep. share (macro)	Modified rep. inc. share (macro)	Exam-corr. inc. share (macro)
P0-10	-2.6	-2.1	-1.7	-1.9	-1.7	-1.9	-1.9
P10-20	1.0	1.0	1.0	1.1	1.0	1.0	1.0
P20-30	2.1	2.1	2.1	2.2	2.2	2.1	2.2
P30-40	3.2	3.4	3.3	3.3	3.4	3.3	3.4
P40-50	4.7	4.8	4.7	4.7	4.9	4.6	4.7
P50-60	6.4	6.5	6.3	6.3	6.5	6.2	6.4
P60-70	8.6	8.7	8.5	8.4	8.7	8.3	8.6
P70-80	11.7	11.6	11.4	11.3	11.6	11.2	11.6
P80-90	16.6	16.4	16.2	16.1	16.4	15.9	16.3
P90-95	12.0	11.8	11.9	11.6	11.9	11.5	11.8
P95-99	16.1	16.0	16.5	16.2	16.0	16.1	16.2
P99-99.5	4.3	4.2	4.4	4.4	4.2	4.4	4.4
P99.5-99.9	6.7	6.5	6.9	6.8	6.5	7.1	6.8
P99.9-P99.95	2.0	1.9	2.0	2.0	1.9	2.1	2.0
P99.95-P99.99	3.2	3.0	2.8	3.2	2.9	3.3	3.1
P99.99-100	4.2	4.1	3.7	4.4	3.9	4.5	4.4
Top 1%	20.3	19.8	19.7	20.7	19.3	21.6	20.6

Note: This Table presents estimates of the distribution of total income, averaging over tax years 2006–2013. In each column, we report the share of total income attributable to different bins of estimated true income. Different columns reflect different specifications. We begin in the first column with reported incomes, ranking individuals by their position in the reported income distribution. We add detected under-reporting only in the next column, estimating the distribution of exam-corrected incomes, ranking individuals by their exam-corrected income. We include these first two columns for comparison purposes; all other columns include detected and undetected under-reporting in total income. Following these two, in the next column we specify undetected under-reporting according to DCE2001 and then estimate total income shares. We report estimates using macro allocation methods in the final four columns.

FIGURE 8: CONTESTED AMOUNTS IN OPERATIONAL AUDITS



Note: This Figure, drawn from GLRRZ, illustrates how the rate at which taxpayers dispute their assessments varies with income, using data on all operational audits of individuals. The top series depicts the share of the total initial tax assessment from audit that is contested by the taxpayer (a dollar-weighted statistic). 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 (an individual-weighted statistic). The data are pooled for fiscal years 2007-2018. The contest rate increases significantly around the 90th percentile and it rises very sharply within the top 10% (up to almost 60% for the top 0.01%). The dollar-weighted contest rate is significantly larger than individual-weighted contest rate, indicating that those with higher assessed values are much more likely to contest their assessments. The large contested shares in the bottom 10% of the distribution are attributable to high-wealth individuals with negative reported incomes; most assessments in this range involve taxpayers claiming large losses that are disallowed on audit.