

THE DYNAMIC EFFECTS OF TAX AUDITS

Arun Advani, William Elming, and Jonathan Shaw*

Abstract—We study the effects of audits on long run compliance behavior using a random audit program covering more than 53,000 tax returns. We find that audits raise reported tax liabilities for five years after audit, effects are longer-lasting for more stable sources of income, and only individuals found to have made errors respond to audit. A total of 60%–65% of revenue from audit comes from the change in reporting behavior. Extending the standard model of rational tax evasion, we show that these results are best explained by information revealed by audits constraining future misreporting. Together these imply that more resources should be devoted to audits, audit targeting should account for reporting responses, and performing audits has additional value beyond merely threatening them.

I. Introduction

AUDITS are a widely used public-policy tool for reducing corruption (Bobonis, Cámara Fuertes, & Schwabe, 2016; Avis, Ferraz, & Finan, 2018), improving public service delivery (Zamboni & Litschig, 2018; Lichand, 2016; Gerardino, Litschig, & Pomeranz, 2020), ensuring environmental standards (Duflo et al., 2013, 2018), and improving tax compliance (Kleven et al., 2011; Pomeranz, 2015; Asatryan & Peichl, 2017; Bergolo et al., 2020; Sarin & Summers, 2020, among others). But audits are costly, so determining how many to do and how best to allocate them are key policy questions (Slemrod & Yitzhaki, 2002). In tax, the standard approach to setting the number of audits is to compare their costs with the expected missing tax uncovered at audit—the static gain from an audit (Allingham & Sandmo, 1972; Kolm, 1973; Yitzhaki, 1987; Bloomquist, 2013). However, audits may change taxpayer behavior. A field experiment in Denmark, which followed taxpayers for a year after audit, found an increased reported liability worth 55% of the audit adjustment (Kleven et al., 2011). This suggests that static gains may understate the total gains from audit. However, without a longer horizon, it is hard to know by how much,

or even whether this effect is reversed in subsequent years, as some lab experiments suggest (Maciejovsky, Kirchler, & Schwarzenberger, 2007; Kastlunger et al., 2009).

This paper studies the long-run effect of tax audits on taxpayer compliance behavior. We combine confidential administrative data on the universe of UK tax filers over thirteen years with a randomised audit programme. We show three main results. First, audits raise subsequent tax reports, but the effect declines to zero over five to eight years. The aggregate additional revenue after audit is at least 1.5 times the underpayment found at audit, implying substantially more resources should be dedicated to audit than a static comparison would suggest. Second, the revenue gain is longer-lasting for more stable income sources. This highlights the importance of dynamics for *targeting* audits, as well as for setting their level. Third, using an event study strategy, we show that these effects are driven by individuals who were found to be underreporting, while there is no response for those found to have reported correctly. These three results can be explained by a model in which audits provide the tax authority with information about a taxpayer's income at the time of audit. This makes later misreporting more difficult, particularly for stable income sources.

To estimate the long-run effect, we exploit a random audit programme run by the UK tax authority (HM Revenue and Customs, HMRC). Over 53,000 individual tax filers were unconditionally randomly selected for audit by the programme between 1998/1999 and 2008/2009, allowing us to address the common concern that audits are typically targeted towards taxpayers believed to be underreporting. Similar to Denmark (Kleven et al., 2011) and in contrast to the United States (Slemrod, Blumenthal, & Christian, 2001; DeBacker et al., 2018; Perez-Truglia & Troiano, 2018), taxpayers are not told these audits are random. This is important as taxpayers may respond differently—likely less—to audits they know are random, relative to when they think the tax authority is concerned about something on their return. We combine these audit data with data on the universe of UK self-assessment taxpayers—individuals who self-file taxes rather than having all tax collected via withholding—from 1998/1999 to 2011/2012. This allows us to follow individuals for many years after audit. For our first identification strategy, we construct a control group for each year of the programme from individuals who could have been selected for a random audit that year but were not. We then study the difference in reporting behavior over time.

Our first result is that *dynamic effects* are positive and substantial: taxpayers report higher levels of tax for five to eight years after audit. We see an initial increase, and then a steady decline, in total tax reported over time. By eight years after audit there is no difference in average tax paid between audited and unaudited taxpayers, though differences are not

Received for publication December 16, 2019. Revision accepted for publication April 8, 2021. Editor: Rema N. Hanna.

*Advani (corresponding author): University of Warwick, CAGE Research Centre, the Institute for Fiscal Studies (IFS), and the Tax Administration Research Centre (TARC); Elming: IFS and TARC at the time of involvement in this work; Shaw: Financial Conduct Authority.

The authors thank Michael Best, Richard Blundell, Tracey Bowler, Monica Costa Dias, Dave Donaldson, Mirko Draca, James Fenske, Clive Fraser, Claus Kreiner, Costas Meghir, Gareth Myles, Matthew Notowidigdo, Aureo de Paula, Andreas Peichl, Imran Rasul, Chris Roth, Joel Slemrod, Hannah Tarrant, and seminar participants at the Tax Systems Conference, Royal Economic Society, Louis-André Gérard-Varet, European Economic Association, Warwick Applied Workshop, OFS Empirical Analysis of Tax Compliance, International Institute of Public Finance, Econometric Society European Meetings, and National Tax Administration Conferences for helpful comments. We also thank Yee Wan Yau and the HMRC Datalab team for assistance with data access. This work contains statistical data from HMRC which is Crown Copyright. The research data sets used may not exactly reproduce HMRC aggregates. The use of HMRC statistical data in this work does not imply the endorsement of HMRC in relation to the interpretation or analysis of the information.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_01101.

statistically significant beyond five years. A total of 60%–65% of the total revenue received as a result of audit comes from this change in reporting behavior. Taking into account this effect, tax authorities should do many more audits: accounting for dynamic effects, even random audits provide a return equal to 80% of their cost to the tax authority. Given the recent focus on the value of audits purely as a threat (Slemrod, Blumenthal, & Christian, 2001; Fellner et al., 2013; Dwenger et al., 2016; Mascagni, 2018; Bergolo et al., 2020; Lichand, 2016); this highlights a benefit of actually performing the audits.

Second, we show that dynamic effects fall to zero slower for more stable income sources. Pension income, which is highly autocorrelated (“stable”) in the absence of audit, responds permanently. At the other extreme, the effect on self-employment and dividend income returns to zero within three years. This is important for two reasons. First, it has implications for the targeting of audits. Going after a smaller suspected discrepancy on a more stable income source can have high returns once dynamic effects are included. Reauditing is also more likely to produce additional yield for individuals with less stable income sources. Second, it is relevant for understanding why people respond to audits, as we describe below. A natural concern in treating this difference causally, and using it to interpret behavior, is that individuals with different types of income may respond differently. We account for this by using pairwise comparisons of income sources within individuals who have both sources, and we demonstrate that the less stable source still declines more quickly.

Third, we show that audits only change the behavior of those who are found to have misreported. To do this we use an event study approach. We compare individuals who were audited at some point in our sample and who ultimately all had the same audit outcome, for example were found to be non-compliant. Allowing for individual and calendar time fixed effects, the comparison is essentially between those whose noncompliance has already been uncovered by a random audit and those who will have it uncovered in the future. We find that being audited only changes the behavior of those who are found to have misreported, and this is true whether or not they received a penalty. Importantly, this tells us that the effect of audits comes not merely from scaring all taxpayers into paying more, but specifically from changing the behavior of those who were previously misreporting. It also allows us to rule out audits reducing tax reports, even for those who were found compliant, in contrast with results using alternative identification strategies (Gemmell & Ratto, 2012; Beer et al., 2020).

These results are consistent with audits providing the tax authority with information at a point in time, which constrains future misreporting. To see this, we extend the canonical model of tax evasion (Allingham & Sandmo, 1972; Yitzhaki, 1987; Kleven et al., 2011) to incorporate (simple) dynamics in the response to audit. This allows us to study the distinct predictions of three different mechanisms that might drive changes in reporting: (i) changes in beliefs about the

underlying audit rate or penalty for evasion (“belief updating”); (ii) changes in the perceived reaudit risk following audit (“reaudit risk”); and (iii) updates to the information held by the tax authority (“information”). Kleven et al. (2011) note that their observed increase in reported tax one year after audit could be explained by some combination of beliefs and reaudit risk, but they cannot disentangle the two. We note that a response to belief updating should be permanent, as taxpayers revise the expected cost of noncompliance (up or down). This is inconsistent with the declining pattern of dynamic effects we see. A response to reaudit risk would decline over time. Whether it took the form of a “bomb crater” (Mittone, 2006)—that the probability of audit is lower in the years following an audit before rising back to baseline—or a worry of higher levels of short-term scrutiny, we should see the same effect across all income sources. We see a positive dynamic effect, ruling out “bomb craters,” and we see a differential decline across income sources, even within individuals, ruling out an effect driven purely by reaudit risk. Instead we propose a third, novel, possibility. As Kleven et al. (2011) note, when taxpayers know the tax authority has access to third-party information about some income source, they are much less likely to underreport. Similarly, when the tax authority performs an audit, it gets a snapshot of income at a point in time. Implausibly large deviations in reported income in following years are likely to trigger an audit, because tax authorities (partly) condition audit selection on differences between reported income and their expectation of that income based on other sources of information (Advani, 2022). As time passes, the snapshot becomes less informative about what current income is likely to be. This is particularly true for less stable income sources. In this case, we should see a decline in dynamic effects over time, with less stable income sources showing a faster decline. We should also only see responses from individuals who were found to have misreported, because no new information about the other taxpayers is revealed to the authority. These are precisely the patterns that are observed.

Our results imply that audits themselves are important, beyond the “fear” or “threat” of audit. Much of the recent literature studying the administration of taxes and the policies that can improve taxpayer compliance has focused on “letter experiments”: how different forms and content of information provided to taxpayers can change their behavior (see Blumenthal, Christian, & Slemrod, 2001; Slemrod et al., 2001 for early work, and Mascagni, 2018; Alm, 2019; Pomeranz and Vila-Belda, 2019; Slemrod, 2019 for recent surveys of this literature). These all aim to change the perceived probability of audit. They have the benefit that they are a very low-cost policy for a tax authority, yet show substantial (short-term) gains. For example, Bergolo et al. (2020) find, in the context of VAT in Uruguay, that firms do not respond to the actual probability of audit when sent letters informing them of this. Instead, firms increase compliance because thinking about the audit scares them into compliance. This raises a question: can high levels of compliance be achieved, while

reducing the number of audits, by directing more resources towards information campaigns? Our results imply that this is harder than previously thought, as much of the gain from audit is the change in behavior it promotes. This response is driven by the information received by the authority through actually conducting the audit. Threat letters do not provide this information benefit. To understand any substitutability with audits, more information is needed on the long-term effects of such letters: for how long do threats raise compliance, and can repeated threats continue to maintain high compliance rates?

In contrast, third-party information is a more direct substitute for audits. Recent work has shown the importance (and limits) of third-party information for improving compliance (Kleven et al., 2011; Pomeranz, 2015; Kleven, Kreiner, & Saez, 2016; Carrillo, Pomeranz, & Singhal, 2017; Slemrod et al., 2017; Naritomi, 2019). Since this directly reduces the information asymmetry between taxpayer and authority, it will also reduce the information value of audits, which drives the dynamic effects. Conversely, for income sources where third-party information can be hard to come by, audits can be a partial alternative to gathering information from other sources. They will not only improve contemporaneous compliance, but also reduce the scope for future noncompliance. This contrasts with work on firms, which finds complementarity between monitoring and enforcement (Almunia & Lopez-Rodriguez, 2018).

We find no evidence of “backfire” effects, where audits reduce compliance. Worries about backfire effects are common across areas of tax policy (Perez-Truglia & Troiano, 2018). In our context they raise the risk that poorly targeted audits may reduce compliance. Gemmell and Ratto (2012) suggest some reduction in tax reported by individuals who are audited and found compliant, relative to individuals not audited. Similar results are found in the United States by Beer et al. (2020) using a matched difference-in-difference approach. Our event study strategy allows for potential differences in unobservable characteristics between compliant and noncompliant individuals, and finds no backfire. The difference in our results, compared to existing work, also suggests that unobservable differences are important in explaining compliance behavior. Since we find no reduction in overall tax paid, it also suggests that lab experimental evidence of bomb crater effects is not reflected in real-world settings (Maciejovsky et al., 2007; Kastlunger et al., 2009), although we note that not all lab experiments find evidence of such effects (Choo, Fonseca, & Myles, 2013).

Finally, we provide a new theoretical mechanism for why audits have the observed effects. Understanding what motivates compliance is a key question for public policy, and there are rich debates on the extent to which moral versus economic calculations drive behavior (Alm, 2019). We focus on the narrower question of why *audits* affect compliance, and we find that information is the key. To do this, we use evidence from random audits to look at both the time path of dynamic effects across income sources and the effects by audit outcome.

Though earlier work has (separately) studied both of these issues, we show how they can be used to understand why audits change behavior.¹ Our results complement those of Bergolo et al. (2020) and Lichand (2016), who find that the *threat* of audit works through fear and belief-updating, respectively. In contrast, receipt of audit works through a change in ability to misreport without being caught, an effect that cannot occur in the absence of actual audit.

The remainder of the paper is organised as follows. Section II outlines the policy context and data sources. Section III provides evidence on who is noncompliant. Section IV shows how audits affect reporting behavior in overall tax, and by different income sources. Section V uses an alternative identification strategy to estimate the impact by audit outcome. Section VI outlines a model of tax evasion with dynamics in the response to audits, to show which mechanisms might rationalise the observed behavior. Section VII concludes.

II. Context and Data

A. The UK Self-Assessment Tax Collection and Enforcement System

In this paper, we focus on individuals who file an income tax self-assessment return in the UK. Over our sample period (1999–2012) this comprised around nine million individuals, one-third of all individual income taxpayers in the UK.² Income tax is the largest of all UK taxes, consistently contributing a quarter of total government receipts over this period. Most sources of income are subject to income tax, including earnings, retirement pensions, income from property, interest on deposits in bank accounts, dividends, and some welfare benefits. Income tax is levied on an individual basis and operates through a system of allowances and bands. Each individual has a personal allowance, which is deducted from total income. The remainder—taxable income—is then subject to a progressive schedule of tax rates. Table 1 shows the share of individuals in our sample reporting nonzero values for each component of income. When we later study income components separately, we focus on those components where at least 5% of the population report nonzero values.

Since incomes covered by self-assessment tend to be harder to verify, there is a significant risk of noncompliance.

¹A number of studies consider dynamic effects for one or two years after audit (Long & Schwartz, 1987; Erard, 1992; Tauchen, Witte, & Beron, 1993; Kleven et al., 2011; Løyland et al., 2019). Concurrently with this study, DeBacker et al. (2018) have a longer (six-year) horizon, and they also consider income stability, albeit with U.S. audits where taxpayers are explicitly told they are random, which Slemrod (2019) notes “would likely trigger different revaluations of how likely a future audit is, and therefore trigger different behavioral changes” (a similar point is made in Kleven et al., 2011). Effects by audit outcome are studied by Gemmell and Ratto (2012) and Beer et al. (2020).

²Filers include self-employed individuals, those with incomes over £100,000 (lower at the start of the sample period), company directors, landlords, and many pensioners. The remainder have all their income tax collected directly via withholding, so are not required to file. Note that UK tax years run across calendar years—we denote tax years using the later year.

TABLE 1.—SHARE OF TAXPAYERS WITH EACH SOURCE OF INCOME

Income component	Proportion
Interest	.587
Employment	.482
Self employment	.375
Dividends	.370
Pensions	.300
Property	.136
Foreign	.048
Trusts and estates	.010
Share schemes	.002
Other	.030

Annual averages for tax years 1998/1999–2008/2009. Includes only control observations, that is, those selected for placebo audit.

Source: Authors' calculations based on HMRC administrative datasets.

As a result, HM Revenue and Customs (HMRC, the UK tax authority) carries out audits each year to deter noncompliance and recover lost revenue. HMRC runs two types of audit: “targeted” (also called “operational”) and “random.” Targeted audits are based on perceived risks of noncompliance. Random audits are unconditionally random from the population, and are used to ensure that all self-assessment taxpayers face a positive probability of being audited, as well as to collect statistical information about the scale of noncompliance and predictors of noncompliance that can be used to implement targeting.

The timeline for the audit process is as follows. The tax year runs from 6th April to 5th April. Shortly after the end of the tax year, HMRC issues a “notice to file” to taxpayers who they believe need to submit a tax return. This is based on information that HMRC held shortly before the end of the tax year. Random audit cases are provisionally selected from the population of individuals issued with a notice to file. The deadline by which taxpayers must submit their tax return is 31 January the following calendar year (e.g., 31 January 2008 for the 2006/2007 tax year). Once returns have been submitted, HMRC deselects some random audit cases (e.g., due to severe illness or death of the taxpayer). At the same time, targeted audits are selected on the basis of the information provided in self-assessment returns and other intelligence. Random audits are selected before targeted audits, and individuals cannot be selected for a targeted audit in the same tax year as a random audit. The list of taxpayers to be audited is passed on to local compliance teams who carry out the audits. Up to and including 2006/2007, audits had to be opened within a year of the 31 January filing deadline, or a year from the actual date of filing for returns filed late. For tax returns relating to 2007/2008 or later, audits had to be opened within a year of the date when the return was filed. Taxpayers subject to an audit are informed when it is opened, but they are not told whether it is a random or targeted audit, in contrast to work done with U.S. random audits (Long & Schwartz, 1987; DeBacker et al., 2018). Even after audit, taxpayers are limited in what they can learn about the audit process because no details of the programme are made public.³ Approximately one-third of taxpayers on the list passed on to local compli-

³Until the publication of this study, even the audit rates were not public information.

ance teams end up not being audited, largely due to resource constraints.⁴

Those who are audited initially receive a letter requesting information to verify what they have reported. If this does not provide all the required information, the taxpayer receives a follow-up phone call, and ultimately in-person visits until the auditor is satisfied.

Where errors are uncovered, individuals are required to pay the additional tax due, and interest. If noncompliance is deemed to be deliberate, the taxpayer might also face an additional penalty of up to 100% of the value of the underpaid tax.

B. Data Sources

We exploit data on income tax self-assessment random audits together with information on income tax returns. This combines a number of different HMRC datasets, linked together on the basis of encrypted taxpayer reference number and tax year.

Audit records for tax years 1998/1999–2008/2009 come from Compliance Quality Initiative (CQI), an operational database that records audits of income tax self-assessment returns. It includes operational information about the audits, such as start and end dates, and audit outcomes: whether noncompliance was found, and the size of any correction, penalties, and interest.


We track individuals before and after the audit using information from tax returns for the years 1998/1999–2011/2012. This comes from two data sets: SA302 and Valid View. The SA302 data set contains information that is sent out to taxpayers summarising their income and tax liability (the SA302 tax calculation form). It is derived from self-assessment returns, which have been put through a tax calculation process. It contains information about total income and tax liability as well as a breakdown into different income sources: employment earnings, self-employment profits, pensions, and so on. For all of these variables, we uprate to 2012 using the Consumer Prices Index (CPI) to account for inflation, and trim the top 1% to avoid outliers having an undue impact on the results.⁵ We supplement these variables with information from Valid View, which provides demographics and filing information (e.g., filing date). Note that we cannot identify actual compliance behavior after the audit: the number of random audit taxpayers that are reaudited is far too small for it to be possible to focus just on them.

An explicit control group of “held out” individuals was not constructed at the time of selection for audit. We therefore draw control individuals from the pool of individuals who actually filed a tax return (i.e., those who appear in SA302). This creates some differences in the filing history between those selected for audit and those who we deem as controls. In a given year, first-time filers may be issued a notice to file

⁴We address the implications for identification in section IVA.

⁵In online appendix C.2 we show our results are robust to alternative levels of trimming.

FIGURE 1.—CHANGE IN THE PROBABILITY OF AUDIT OVER TIME



Constructed using data on individuals who received an audit of their self-assessment tax return for a tax year between 1998/1999 and 2008/2009, and the full sample of self-assessment returns for the same period.
 Source: Calculations based on HMRC administrative data sets.

after selection for audit has taken place. They may also end up back-filing one or two returns. Since we cannot directly observe the first year in which a notice to file was issued, in our empirical strategy it is necessary for us to control for the length of time each taxpayer has been in self-assessment. More details—including tests to demonstrate this ensures samples are balanced—are given in section IVA below.

III. Tax Evasion in the UK

In this section, we first provide some descriptives on the probability and timeline of audits. We then show that there is significant noncompliance among individual self-assessment taxpayers, both in the share of taxpayers who are found noncompliant and the share of tax that is misreported. More than one-third of self-assessment taxpayers are found to be noncompliant, equal to 12% of all income taxpayers.

A. Audit Descriptives

Figure 1 shows the share of individuals per year who face an income tax random audit over the period 1998/1999–2008/2009. On average over the period, the probabilities of being audited are 0.04% (4 in 10,000) for random audits and 2.8% for targeted audits.

Table A1 provides some summary statistics for lags in, and durations of, the audit process among random audit cases. As described above, up to and including the 2006/2007 return, HMRC had to begin an audit within 12 months of the 31 January filing deadline; since then, HMRC has had to begin an audit within 12 months of the filing date. The average lag between when the tax return was filed and when the random audit was started is 8.9 months, but 10% have a lag of 14 months or more. The average duration of audits is 5.3 months,

TABLE 2.—RANDOM AUDIT OUTCOMES

	Mean	Std. dev.
Proportion of audited returns deemed		
Correct	.532	.499
Incorrect but no underpayment	.111	.314
Incorrect with underpayment (noncompliant)	.357	.479
Mean additional tax if noncompliant (£)	2,314	7,758
Distribution of additional tax if noncompliant		
Share £1–100	.116	.320
Share £101–1,000	.483	.500
Share £1,001–10,000	.361	.480
Share £10,001+	.039	.194
Observations	34,630	

Annual averages for tax years 1998/1999–2008/2009. Includes all individuals with a completed random audit.

Source: Authors' calculations based on HMRC administrative data sets.

but 10% experience a duration of 13 months or more. Taken together, this means that the average time between a return being filed and an audit being concluded is 14.3 months, but there are some taxpayers for whom the experience is much more drawn out: for almost 10% it is two years or more. This means that individuals will generally have filed at least one subsequent tax return before the outcome of the audit is clear, and some will have filed two tax returns. This will be relevant for interpreting the results in section IV.

B. Evidence of Noncompliance

We begin by studying the direct results of random audits, using data on 34,630 completed random audits of individual self-assessment taxpayers from 1998/1999 to 2008/2009.⁶ Table 2 summarises the outcomes of these random audits. More than half of all returns are found to be correct, 11%

⁶53,400 cases were selected for audit over the period, of which 35,630 were implemented.

are found to be incorrect but with no underpayment of tax, and 36% are “noncompliant,” that is, incorrect and have a tax underpayment.⁷ Whilst this is a much higher rate of noncompliance than has been found in other developed country contexts, it should be noted that the self-assessment tax population is a selected subset of all taxpayers. In particular, it covers those for whom a simple withholding of income at source is not sufficient to collect the correct tax. This may be either because some income cannot be withheld (e.g., property or self-employed income), or because PAYE struggles to assign the correct withholding codes (e.g., for people with multiple sources of pension income). Despite this, since self-assessment taxpayers make up a third of all UK taxpayers, this implies an overall noncompliance rate of 8%–12% among all taxpayers.⁸

Turning to the intensive margin, the average additional tax owed among the noncompliant is £2,314, or 32% of average liabilities. Since just over a third of random audits find evidence of noncompliance, the average additional tax owed from an audit is then £826.⁹ However, the distribution is heavily skewed: 60% of noncompliant individuals owe additional tax of £1,000 or less, whilst 4% owe more than £10,000. In terms of total revenue, those owing £1,000 or less make up only 9% of the underreported revenue; the 4% owing more than £10,000 collectively owe more than 42% of the revenue. Equity concerns around noncompliance are well-known: it is seen as unfair that some are not “paying their fair share.” But this variation in noncompliance is also important for economic efficiency. Noncompliant individuals previously acted as though there was a lower tax rate. This makes their activities seem relatively more productive than those of compliant individuals, so it can lead to resource misallocation.

IV. Dynamic Impacts of Audits

In this section we establish two main results. First, we show that audits lead to an increase in reported incomes and taxes in subsequent years. Looking at total income and total tax, this increase lasts five to eight years after the tax year for which the audit was done. Second, we show variation in this impact by income source. In particular, more autocorrelated income sources (such as pensions) seem to respond permanently to audit. In contrast, income sources that are less autocorrelated, such as self-employment income, more quickly return

⁷Incorrect with no underpayment includes those who, for example, owed no taxes because they had legitimate losses, but had overstated those losses so would owe less in future years. Anecdotally, it also includes some cases where actual overpayments of tax were made, although we cannot separately identify which.

⁸This is a lower bound, since it assumes everyone who should be in self-assessment does register, all noncompliance is picked up at audit, and those who do not need to register are also fully compliant. The range from 8% to 12% depends on the assumptions made about the implementation of audits. If, among those selected for audit, implementation of audit were random, this would imply a 12% noncompliance rate. On the other hand, if there is perfect compliance among those for whom audits were not implemented, this would imply an 8% rate.

⁹This is the additional tax owed. A further £101 is owed, on average, in penalties. This is highly concentrated, with less than 7% of those audited owing any penalty amount.

to baseline. This second result will later help explain *why* we see these dynamic responses. Before describing these results in detail, we first discuss the empirical approach taken. Briefly, we compare individuals selected for random audit with those not selected but who could have been selected. We control for filing history to account for the way the sample was selected.

A. Estimation

To understand how audits affect future tax receipts, we want to estimate the change in tax paid in the years after audit that is caused by the audit. We recover this using the “random audits program” run by the tax authority (HMRC). This programme selects for audit a random sample of taxpayers from the pool of taxpayers known to be required to file for a given tax year. One can therefore compare those selected for audit with others who were not selected but who could have been.

In each audited tax year we select a sample of individuals who were not audited and could have been. We assign them a “placebo audit” for that tax year. We can then compare them over time to individuals actually selected for audit for that year. Our sample, therefore, consists of individuals who were selected for random audit in some year between 1999 and 2009, and individuals who could have been selected in those same years but were not. Our data on tax returns go up to 2012. For every individual selected for audit in a given tax year, we draw six control individuals from the population of those who could have been audited in the same tax year.¹⁰

In practice, a little more than two-thirds of those selected for random audit are actually audited. This is explained by the high workload faced by the compliance teams implementing audits. Additionally, a small fraction of the control group (around 2%) is also audited. Random audits are selected before targeted audits, and no explicit control group was constructed to “hold out” some individuals from targeting. To our knowledge, in prior work only Kleven et al. (2011) have an explicit control group. This explains why they can only study a single year after audit—tax authorities are unwilling to hold off on high-value audits for multiple years. Hence we compare those selected for a random audit to a “business as usual” group, rather than a pure control group. This will tend to reduce the estimated impacts, since individuals in the control group who are most likely to be noncompliant are audited.

In the empirical work to follow, we focus on the local average treatment effect (LATE), instrumenting receipt of audit with selection for random audit. This is the relevant number for a tax authority thinking about simultaneously expanding the size of the random audit programme and the number of auditors. It gives the average impact h years after audit for an

¹⁰In principle, the entire population of taxpayers who could have been audited could have been used. However, because the data could be accessed only in a secure facility at the tax office, computational constraints given the available hardware limited the sample size that could be used.

TABLE 3.—SAMPLE BALANCE, CONDITIONING ON FILING HISTORY

Years after audit		−5	−4	−3	−2	−1
Characteristics						
Female	Mean	.274	.276	.278	.282	.287
	Difference	−.005	−.006	−.005	−.006	−.005
	<i>p</i> -value	.236	.212	.292	.234	.338
Age	Mean	49.2	49.3	49.3	49.4	49.5
	Difference	.0	.0	.1	.1	.1
	<i>p</i> -value	.472	.600	.188	.170	.110
In London or SE	Mean	.333	.334	.335	.333	.331
	Difference	−.003	.001*	.003	.002	.002
	<i>p</i> -value	.159	.026	.015	.317	.190
Has tax agent	Mean	.628	.614	.603	.589	.573
	Difference	−.003	−.001	−.001	.002	.002
	<i>p</i> -value	.522	.500	.376	.675	.606
Income and tax totals						
Total taxable income	Mean	35,075	34,670	34,030	32,912	31,755
	Difference	−2	35	−163	71	56
	<i>p</i> -value	.979	.469	.012	.280	.439
Total tax	Mean	9,646	9,539	9,321	8,979	8,635
	Difference	14	12	−40	12	15
	<i>p</i> -value	.982	.288	.061	.261	.887
Income components						
Employment	Mean	22,508	22,534	22,266	21,708	21,145
	Difference	11	−57	−98	112*	43*
	<i>p</i> -value	.758	.023	.152	.049	.05
Self-employment	Mean	6,546	6,379	6,200	5,950	5,581
	Difference	56	38	−49	−18	29
	<i>p</i> -value	.298	.435	.033	.161	.684
Interest and dividends	Mean	4,007	3,905	3,895	3,759	3,645
	Difference	−26	16	−27	7	4
	<i>p</i> -value	.667	.189	.235	.958	.086
Pensions	Mean	3,493	3,542	3,561	3,562	3,531
	Difference	−23	−23	−3	4	22
	<i>p</i> -value	.806	.482	.681	.463	.523
Property	Mean	869	844	811	769	726
	Difference	−5	−6	0	6	0
	<i>p</i> -value	.282	.209	.525	.072	.518

“Years after audit” measures time relative to audit, or placebo audit for controls. “Mean” is the mean outcome in the control (not selected for audit) group across all years. “Difference” is the coefficient on the treatment dummy in a regression of the outcome on a treatment dummy and dummies for whether the taxpayer filed taxes in each of the four years before audit (or placebo audit for controls). Treatment dummy equals 1 if taxpayer was selected by HMRC for a random audit. *p*-values are derived from an F-test that coefficients on interactions between treatment and tax year dummies are all zero in a regression of the outcome of interest on tax year dummies, interactions between treatment and tax year dummies, and dummies for whether the taxpayer filed taxes in each of the four years before audit (or placebo audit for controls). This is a stronger test than just testing the coefficient on treatment not interacted. Monetary values are in 2012 prices. Standard errors are clustered by taxpayer. * *p* < .05, ** *p* < .01, and *** *p* < .001.

Source: Authors’ calculations based on HMRC administrative data sets.

additional random audit case that might be worked, against which the cost of the audit would be compared.

One limitation of our data is a slight mismatch between our treated and control samples in terms of their probability of filing in previous years, for reasons relating to the audit timeline and when they were first issued a notice to file, as described in section IIB. This can be seen in table A2, which documents (unconditional) sample balance between five and one years before audit, for income and tax totals, income components, and individual characteristics. Overall balancing statistics suggest that the samples are fairly well-balanced: the *p*-value of the likelihood-ratio test of the joint insignificance of all the regressors is 0.181, while the mean and median absolute standardised percentage bias across all outcomes of interest are low at 2.4% and 1.7%, respectively.¹¹ However,


the likelihood of being in the sample in previous years (“survival”) differs between our treatment and control groups. This difference is consistent with how the treatment and control groups were selected, so it might reflect real differences in the samples. We therefore include controls for presence in the data in the years before audit.¹² Table 3 shows that once we condition on past survival, the sample is balanced.

the average of the sample variances in the treated and control groups (see Rosenbaum & Rubin, 1985). Rubin’s B and R statistics are also well within reasonable thresholds to consider the samples to be balanced, at 10.8 and 0.983, respectively. Rubin’s B is the absolute standardised difference of the means of the linear index of the propensity score in the treated and control group. Rubin’s R is the ratio of treated to control variances of the propensity score index. Rubin (2001) recommends that B be less than 25 and that R be between 0.5 and 2 for the samples to be considered sufficiently balanced.

¹²In online appendix C.1, we show the results taking a different approach, where we instead use stratified random sampling conditioning the stratification on filing history. Point estimates are similar, and never statistically significantly different from our main approach, although they decline more rapidly from year four.

¹¹The standardised percentage bias is the difference in the sample means between treated and control groups as a percentage of the square root of

FIGURE 2.—DYNAMIC EFFECT OF AUDITS ON TOTAL REPORTED TAX OWED



Sample includes individuals selected for a random audit between 1998/1999 and 2008/2009, and control individuals who could have been selected in the same years but were not. It uses tax returns from 1998/1999 to 2011/2012. The solid line plots the point estimate for the difference in average “total reported tax” between individuals who were and weren’t audited, for different numbers of years after the audit. This comes from a regression of total reported tax on dummies for years since audit (or placebo audit for controls), dummies for years since audit (or placebo audit for controls) interacted with treatment status, tax year dummies, and dummies for whether the taxpayer filed a return in each of the four years before audit, with audit status instrumented by selection for audit. Standard errors are clustered at the individual level. Source: Calculations based on HMRC administrative data sets.

We therefore estimate the following specification:

$$Y_{ihts} = \sum_{h=-5}^8 \alpha_h \eta_h + \sum_{h=-5}^8 \beta_h \eta_h D_i + \sum_{s=1999}^{2012} \gamma_s T_s + \sum_{s=-4}^{-1} \delta_s S_{is} + \varepsilon_{ihts}, \tag{1}$$

where Y_{ihts} is the outcome for individual i , h years after the tax year selected for audit (with control observations having $h = 0$ for the tax year for which they were drawn as controls), when the current calendar year is $s \equiv t + h$. η_h are indicators for being h years after the tax year selected for audit; D_i is an indicator for whether the individual is actually audited; T_s is a calendar time indicator for tax year s ; and $\{S_{i,-1}, \dots, S_{i,-4}\}$ are indicators for whether the individual was in the data in each of the four years before audit. The error term, ε_{ihts} , is clustered at the individual level. Audit status, D_i , is instrumented by (random) selection for audit, Z_i . The coefficients of interest are $\beta_h \forall h$. These estimate the impact of the audit on the outcome variable h years after the tax year selected for audit, measured as the difference in the mean outcome for those actually audited and those who would have been audited only if selected for a random audit.

B. Overall Impact of Audits

Beyond the direct effects of the audit, described in section II, we also see clear evidence of *dynamic* effects. Comparing individuals who were randomly selected for audit with individuals who could have been (but were not) selected, those selected for audit on average report higher levels of tax owed in the years after audit. Figure 2 shows the estimated

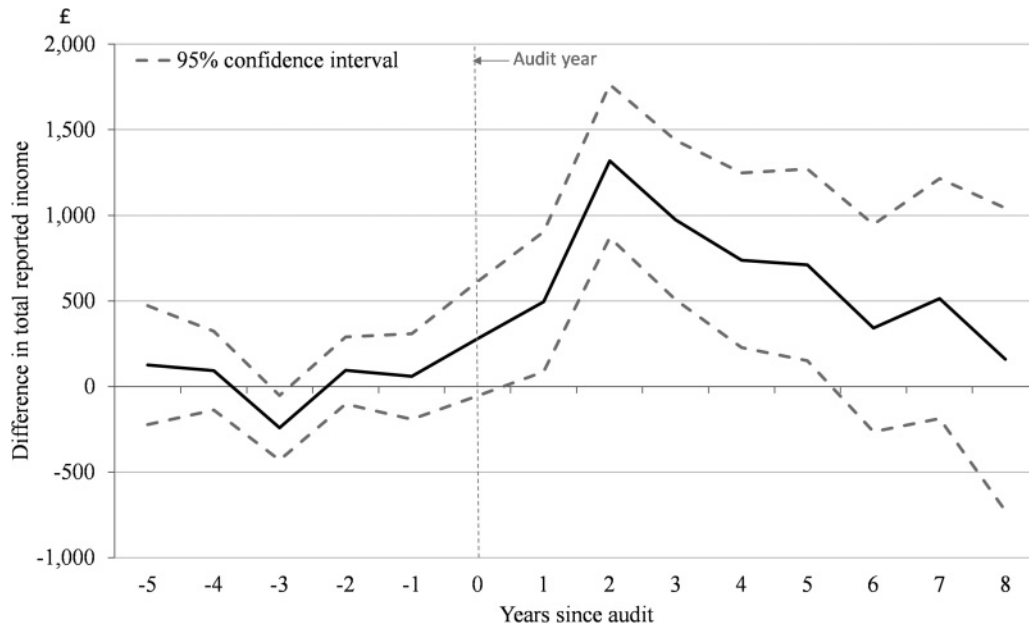
impact on those who were actually audited (i.e., the LATE). The difference in the share audited between the treated and control group is around 66 percentage points, so the LATE is around 1.5 times the intention to treat estimate.

The impact of an audit peaks two years after the tax year for which the audit is conducted. This is consistent with the fact that many audits are not started until after the following year’s tax return has already been submitted.¹³ Reported tax among audited taxpayers is significantly greater than among nonaudited taxpayers for five years after the audit, and the point estimate appears to decline relatively smoothly, getting close to zero by the eighth tax year after the audited year. This pattern of effects is robust to changes in the level of trimming, although, when lower levels of trimming are used, standard errors are larger and consequently some significance levels are lower (see online appendix C.2 for details).

From figure 2, we can estimate how much revenue audits raise on average by changing the behavior of audited individuals. Over the five (eight) years after the audited year, the dynamic effects bring in an additional £1,230 (£1,530), 1.5 (1.8) times the direct effect of audit. Although taxpayers in the United States are explicitly told that the random audits

¹³In our sample, almost a quarter of audits are not opened for more than 12 months from the date of filing (see table A1). Additionally, there can be some lag between the tax authority “taking up” a case for audit and notification being received by the taxpayer. If taxpayers each consistently file at the same time every year, this implies at least one-quarter would have filed without knowledge of the audit. More than half will have filed without knowing the result of the audit (table A1). One could instead set $h = 0$ as the time at which audit begins, but this information is not available for controls, so it risks creating bias if the timing of opening audits among individuals selected for audit is nonrandom.

FIGURE 3.—DYNAMIC EFFECT OF AUDITS ON TOTAL REPORTED INCOME



Sample includes individuals selected for a random audit between 1998/1999 and 2008/2009, and control individuals who could have been selected in the same years but were not. It uses tax returns from 1998/1999 to 2011/2012. The solid line plots the point estimate for the difference in average “total reported income” (income from all sources) between individuals who were and weren’t audited, for different numbers of years after the audit. This comes from a regression of total reported income on dummies for years since audit (or placebo audit for controls), dummies for years since audit (or placebo audit for controls) interacted with treatment status, tax year dummies, and dummies for whether the taxpayer filed a return in each of the four years before audit, with audit status instrumented by selection for audit. Standard errors are clustered at the individual level.

Source: Calculations based on HMRC administrative data sets.

are random, DeBacker et al. (2018) find a similar ratio between direct and indirect effects of audit. Ex ante one might have expected smaller behavioral effects, because taxpayers are aware that the authority is not acting based on any suspicion of wrongdoing. Our exploration of the mechanism driving these dynamics will explain why, ex post, these effects should be so similar: the dynamics are driven by constraints to misreporting caused by audit, rather than belief-updating or perceived reaudit risk, both of which may respond to the reasoning behind the audit.

These dynamic effects highlight the policy importance of studying the long-term impact of audits: when determining the audit strategy, the revenue-raising effects of audits would be grossly understated without considering the impact on future behavior. This would imply too few audits taking place.

It is important to note that the optimal number of audits will in general not equate the marginal return on audit to the marginal cost of an audit. Audits require real resource costs, while the direct benefits are a transfer of resources from citizens to the state (see Slemrod & Yitzhaki, 1987 for a longer discussion of this point). There are likely also indirect benefits in terms of maintaining overall compliance, as well as potentially intrinsic value placed in upholding the rule of law (Cowell, 1990). Additionally, the social cost of audit must incorporate not only the cost to the tax authority, but also the cost to the taxpayer for which accurate figures are difficult to come by (Burgherr, 2021). We therefore do not attempt a full welfare analysis. Instead we merely note that dynamic effects increase the resources that are transferred to the state without increasing the administrative costs of audit.

Assuming that a positive weight is placed on such transfers, taking into account dynamic effects increases the number of audits that should be undertaken.

Figure 3 shows that a very similar pattern holds for the impact on *total income* reported. Again there is a clear dynamic effect, peaking two years after the audited year and declining to zero by year eight, though not significantly different from zero by year five. This provides additional support to the previous result for tax, and is not purely by construction, because expenses can often be used to offset income to reduce tax (Carrillo et al., 2017; Slemrod et al., 2017).

C. Impact by Income Source


We repeat the previous estimation separately by income sources, focusing on income sources for which at least 5% of the sample report nonzero amounts.¹⁴ This will be one way in which we discriminate between different possible explanations for why we see dynamic effects.

Figure 4 shows how the impact of an audit changes over time for the different components of income. Since the magnitudes of these incomes are different, for comparability we rescale them relative to the peak impact for that income source.

We see that, relative to the peak, self-employment income and dividends decline relatively quickly. Three years later point estimates for these are close to zero, that is, reporting is

¹⁴We exclude interest income, because it is very small and not everyone needs to report this, making it hard to compare. See table 1 for information on the share of individuals with each income source.

FIGURE 4.—RELATIVE DYNAMICS BY INCOME SOURCE: LESS AUTOCORRELATED SOURCES OF INCOME SEE FASTER DECLINES



Sample includes individuals selected for a random audit between 1998/1999 and 2008/2009, and control individuals who could have been selected in the same years but were not. It uses tax returns from 1998/1999 to 2011/2012. Each line plots the point estimate for the difference in the average of a particular component of income between individuals who were and weren't audited, for different numbers of years after the peak impact for that income source. This comes from a regression of each income component on dummies for years since audit (or placebo audit for controls), dummies for years since audit (or placebo audit for controls) interacted with treatment status, tax year dummies, and dummies for whether the taxpayer filed a return in each of the four years before audit, with audit status instrumented by selection for audit.

Source: Calculations based on HMRC administrative data sets.

TABLE 4.—AUTOCORRELATION BY INCOME SOURCE

	Corr($t, t - 1$)	Corr($t, t - 2$)	Corr($t, t - 3$)
Pension income	.946	.904	.864
Property income	.896	.836	.790
Employment income	.862	.769	.690
Interest income	.835	.722	.640
Self-employment income	.832	.728	.644
Dividend income	.813	.723	.657
Observations	4,506,548	4,506,548	4,506,548

Annual averages for years 1998/1999–2011/2012.

Source: Calculations based on HMRC administrative data sets.

not different to the control group. In contrast, pension income exhibits little decline. Six years later it retains 80% of the impact, and this is not statistically different from 100%. This pattern is suggestive of the importance of autocorrelation: income sources that one would expect to be more correlated over time appear to show weaker declines.

Table 4 shows the autocorrelation for each income source. Pension income is highly autocorrelated because it will typically be an annuity and therefore fixed over time; property income is slightly less stable because rents may vary more; and at the other extreme, self-employment and dividend income are considerably less stable. The relative autocorrelations of income sources line up exactly with their speeds of decline.¹⁵

There are two caveats to these results. The first is that these measures are noisy, so if confidence intervals were added to

¹⁵Note that a comparison of pensions versus property income is helpful in distinguishing this effect of autocorrelation compared with the effect of third-party information. Both have a high autocorrelation, but pension income was third-party reported while property income was not. In figure 4 we see essentially the same effect for both sources, despite the large difference in third-party information. Conversely, comparing property income and dividend income—which, like property, is also not third-party reported but has a low autocorrelation—we see very different effects.

figure 4 for each income source, many would overlap. The second is that individuals with different income sources may have different propensities for noncompliance.

To tackle these concerns, we next use two alternative strategies. First, we compare within individuals who have multiple income sources. This immediately solves the second problem above because our results will be within individuals. It will also lead to ten pairwise comparisons: every unordered pair of the five income sources studied. For each pair, our sample is composed of individuals who had both sources sometime in the three years before audit. We then study the relative fall in reporting of each of these income sources four years after the peak. In each case, we expect to find that the less autocorrelated source falls fastest.

We find this result in eight out of ten cases. If there were no relationship, we should find this to be true in around five of the tests. The probability of this result under the null of no relationship is 5.5%, close to standard significance thresholds. Hence more autocorrelated income sources do seem to decline more slowly than less autocorrelated ones.

Our second strategy to tackle concern about heterogeneity in who receives different income sources is to reweight individuals based on individual characteristics. This ensures that the distribution of observed characteristics is the same across recipients of different incomes. We divide individuals into groups by sex, age band (below 40, 40–65, and above 65—the UK state pension age at which people typically retire), and quartiles of filing history. We then run weighted regressions so that the weighted samples match closely the distribution of these characteristics seen among individuals with self-employment income. We replicate figure 4 using the results of the reweighted regression, shown as figure A2. The results look very similar—the only noticeable effects are

that property income appears to decline slightly faster than previously, and dividend income much faster.

Our interpretation for this result, which we formalise below, is that audits provide the tax authority with information. Where errors are uncovered, taxpayers file amended returns. Although we do not know, and would not be allowed to reveal, precisely how audit targeting is done, it is clear that “surprising” deviations from recorded historic reports are part of this. The amended return is therefore creating a new benchmark against which future returns will be compared. Hence, income from highly autocorrelated sources will—once uncovered—be hard to hide again, as deviations from the truth will be easily noticed. In contrast, declines in less autocorrelated income sources are less informative to the authority because they may well be real for an individual taxpayer. Viewed in aggregate, falls and rises should be equally likely, because the control group will account for any trends in the income source. Hence when we observe a decline in aggregate income reports (e.g., for dividend income among audited taxpayers), this can be attributed to noncompliance, although we cannot identify which individuals are the ones underreporting. Because declines are faster for less autocorrelated income sources, this suggests the importance of information provision. This is something we know to be important from other settings (Kleven et al., 2011; Pomeranz, 2015), although the value of audits as a potential source of information about future tax has not previously been recognised.

One caveat to this interpretation is that falls in reporting could alternatively be driven by changes in actual income. For example, those who are audited might sell shares to pay fines, reducing dividend income. Whilst this is possible, it seems unlikely. In cash terms, the peak additional income reported for those who have dividend income is £414. Assuming a high-end estimate for the dividend yield of 10% implies £4,140 of undeclared shares. Conservatively assuming also that individuals are on the higher rate of income tax, this implies an additional £135 of tax owed. The absolute maximum penalty for misreporting is 100% of the tax due (on top of paying the tax). So selling all these shares (and hence looking like the control group) would be needed only for an individual who is found to have misreported for at least fifteen years, and receives the maximum fine. While such cases might exist, it seems extreme to assume that this is occurring on average. Hence we think it is unlikely that the observed pattern represents changes in real behavior, rather than reporting, though we cannot definitively rule it out.

V. Impacts by Audit Outcome

We next consider how dynamic effects vary depending on the outcome of audit. This is important for policy, as it helps distinguish whether merely the process of being audited is enough to impact reported income and tax. We find that those who were found to be correct do not respond, while those for whom errors were found increase reported tax. Being audited

per se does not appear to increase reported tax—that is, there is no change in behavior among compliant taxpayers—but those found to have underpaid are 18 percentage points more likely to report higher tax owed after audit. We first describe the approach taken to study this question, because our previous control group cannot help us study effects by audit outcome. We then describe the findings highlighted above.

A. Empirical Approach

Since we now wish to study audit impacts separately by audit outcome, we cannot use the earlier identification strategy. In the “placebo audit” group, we cannot observe what audit outcomes would have been, so we cannot construct separate control groups for each audit outcome. Gemmell and Ratto (2012) studied this question by comparing each treatment group to the original control group containing people with a mix of possible outcomes, implicitly assuming that audit outcomes are exogenously assigned. More recently, Beer et al. (2020) used a matched difference-in-difference approach, allowing for observable differences in audit outcome.

We take an event study approach to answer this question. Our sample for each regression is the set of observations for individuals who are audited and found to have some particular outcome (e.g., found to be compliant). Within that sample, the timing of audit is random—there is nothing systematic that led individuals to be selected in a particular year within the sample. Hence we can compare the outcome for someone audited and found to have a particular status (e.g., to be compliant) with someone who *will be* audited and found to have the same status.

For our variable of interest, we now focus on a binary variable measuring whether tax paid increases, rather than on the sizes of the increase, as in Pomeranz (2015). In particular, we estimate a linear probability model in which the outcome is whether tax paid in year t is larger than in the year before audit. Our interest now is understanding *which* individuals—when split by audit status—respond. This outcome is therefore preferred because it compares individuals to their own history, and it is equally responsive to increases for individuals across the distribution of taxes owed. It is also less sensitive to relatively extreme observations, which is more important in our event study approach because the sample size is now much smaller. Whereas previously we had a treatment group of 53,000 individuals, and could draw a large sample of controls from the nonaudit population, now the entire sample is those selected for audit. That sample is then further split into subsamples by audit outcome status, making results more sensitive to outliers and reducing power. Use of a binary variable removes this sensitivity without limiting our ability to study which groups respond.

In our specification we control for a number of key covariates: sex, age, industry, region, and years filing, as well as calendar-year fixed effects. Many of these individual characteristics have been shown to be predictive of noncompliance (Advani, 2022), so if responsiveness to audit also differs

TABLE 5.—IMPACT BY AUDIT OUTCOME

Years since audit	Overall	Correct	Mistake non-positive	Mistake positive	Positive yield + penalty	Not audited
-5	-.006 (.013)	-.042* (.018)	.048 (.049)	.033 (.032)	-.014 (.072)	-.002 (.030)
-4	.007 (.014)	-.034 (.019)	.068 (.049)	.050 (.033)	.037 (.068)	-.006 (.030)
-3	.005 (.014)	-.023 (.019)	.058 (.050)	.039 (.033)	.042 (.068)	-.016 (.030)
-2	.022 (.014)	-.005 (.019)	.079 (.050)	.075* (.033)	.032 (.068)	-.008 (.030)
-1	<i>Outcome is difference from -1 so zero by construction</i>					
0	.056*** (.014)	.016 (.019)	.131* (.051)	.179*** (.033)	.092 (.069)	-.014 (.030)
1	.048*** (.014)	.012 (.019)	.109* (.051)	.174*** (.033)	.180** (.069)	-.037 (.030)
2	.042** (.013)	.007 (.020)	.135** (.051)	.152*** (.033)	.207** (.069)	-.052 (.031)
3	.030* (.014)	-.007 (.020)	.135** (.052)	.133*** (.034)	.171* (.069)	-.048 (.031)
4	.031* (.014)	-.0024 (.021)	.134* (.052)	.137*** (.034)	.143* (.070)	-.045 (.031)
5	.033* (.016)	.019 (.023)	.160** (.056)	.119** (.037)	.128 (.074)	-.052 (.034)
Observations	124,223	46,911	9,519	25,666	6,983	35,144

The outcome variable is a dummy for whether tax paid is higher in each of the years before/after audit than the year immediately before audit ('-1'). "Overall" uses the full sample of audited individuals to perform an event study for whether tax paid is higher than in the year before audit. Coefficients from a linear probability model are shown, with standard errors in parentheses. Other columns split the audited sample by audit outcome: tax return found to be correct; tax return found to have a mistake but which doesn't change tax liability (or in a small number of cases reduced liability); tax return found to have a mistake leading to increased tax liability, but no penalty charged (i.e., treated as legitimate error); tax return found to have underreported liability and a penalty charged (i.e., deemed to be deliberate); tax return selected for audit but no audit actually implemented (placebo test). * $p < .05$, ** $p < .01$, and *** $p < .001$.

Source: Authors' calculations based on HMRC administrative data sets.

by these characteristics, then without such controls we may partly pick up a purely compositional effect.

B. Results by Audit Outcome

To assess the reasonableness of the approach, we begin again by studying the estimated impact in the years before audit. The first four rows of table 5 provide the results for the preaudit period. It can be seen that all the point estimates are close to zero, providing support for the validity of this approach. A second test of validity can be seen from the "Not audited" column. This estimates the effect of being selected for audit on individuals who were never actually audited, nor informed that they had been selected. As expected, again the point estimates are very close to zero.

Turning to the other columns, three results can be seen. First, those who were audited and found to have made no errors do not respond. This is important because it tells us that the dynamic response isn't driven by the mere fact of audit. Direct audit effects could happen, for example, if the process of audit were sufficiently unpleasant that taxpayers decided to err upwards when uncertain in the hope of avoiding further audits. One could also potentially have seen negative direct effects in this group. If some taxpayers were incorrectly found to be compliant, they may learn that the tax authority is less effective at detecting noncompliance than they previously believed, and reduce payments. We find neither of these results: on average, those whose returns are found correct do not change their reports, in contrast to work by Gemmill and Ratto (2012) and Beer et al. (2020).

Second, those who are found to have made errors are more likely to report higher levels of tax in subsequent years. Even four years later they are 13–14 percentage points more likely to report higher tax owed. Hence the long-term effects observed appear to all come from correcting errors made by the taxpayer. Note that even those who made errors but owed no additional tax respond to the audit. This is because the errors made might affect future tax liability. For example, claiming excessively large expenses today might increase the size of a loss on property income that can be carried forward: correcting this increases future tax liabilities. Anecdotally, from speaking to audit officers, in some cases these individuals shift their reports to pay tax in the audit year so that they can smooth out the additional tax liability that they will now face over the coming years.

Third, those who receive a penalty appear to have been driving some of the shape of the dynamics we observed earlier, where we saw a peak two years after the year selected for audit. Whilst those with mistakes but no penalty respond immediately, the response for those with a penalty peaks two years after the year for which the audit is done. This reflects two features of the audit process. First, those who ultimately receive penalties typically take longest to audit, because their underreporting requires more work to detect. The audit settlement date is thus later. If some taxpayers wait until the audit (and uncertainty about detection) is resolved to respond, this will delay the time until they are observed to respond. Second, taxpayers with mistakes but no penalties will have their *original* return corrected, so an immediate response is observed. On the other hand, those who receive a penalty

may not have their return corrected: in most cases they instead file a separate form detailing additional tax, interest, and penalties.

Among individual characteristics, the only one which predicts responsiveness overall is sex: women are around 3 percentage points more likely to respond to an audit. This is purely driven by compositional effects. Judging by audit outcome, there are no differences in responsiveness by sex.

VI. Simple Model of Tax Evasion and Audit Response

To help understand the mechanism underlying the observed results, we consider an extended version of the model of rational tax evasion by Allingham and Sandmo (1972), which is based on the Becker (1968) model of crime. In the Allingham and Sandmo (1972) model, individuals receive income and choose how much to report to the authority. Underreporting has the benefit that individuals end up paying less tax, but the cost that they may be caught and receive a punishment on top of paying the correct tax. The probability of being caught is increasing in the amount of evasion. Kleven et al. (2011) extend this to allow some income to be third-party-reported: underreporting this income is detected with probability 1, so individuals will only evade out of non-third-party reported income.

The key innovation of our model is to split non-third-party reported income into more versus less stable sources.¹⁶ Incomes from some sources, such as pension annuity income, are very autocorrelated (“stable”), while other sources, such as self-employment income for a sole trader, are much less stable. Autocorrelation captures the extent to which information learned in an audit today is informative about incomes tomorrow. By first extending the model of Kleven et al. (2011) to multiple time periods, and then allowing for differential autocorrelation of income sources, we are able to distinguish different possible mechanisms for *why* audits are observed to have long-term effects.

Consider an individual who is audited (for the first time) in year t . Being audited may change his/her reporting for some combination of the following three reasons: (i) beliefs about the underlying audit rate or penalty for evasion (“belief updating”); (ii) changes in the perceived reaudit risk following audit (“reaudit risk”); and (iii) updates to the information held by the tax authority (“information”).¹⁷

In the first of these mechanisms, there is a change in beliefs about fixed parameters, either audit rate or penalty. Consequently, any response should also be permanent and common across all income sources. Empirically neither of these is true.

Under the second mechanism, the individual perceives a temporary change in the risk of being audited. If s/he perceives the risk to have risen, s/he should be more compliant in the short term, but as perceived risk returns to baseline, re-

porting should do so as well. Conversely, if s/he perceived the risk to have fallen—the so-called “bomb crater effect” (Mittone, 2006; Maciejovsky et al., 2007; and Kastlunger et al., 2009)—then s/he should be temporarily less compliant. In both cases, the dynamics of this behavior should be common across income sources. The differential responses across income sources, even within individuals, are not consistent with this mechanism.

The final mechanism is that audits provide information that differentially changes the ability to hide certain sources of income. Performing an audit provides the tax authority with more accurate information on a taxpayer’s income at a point in time. In subsequent years, information from the audit will make evasion of more stable income sources easier to detect, but for less stable income sources the effect will rapidly wear off. Hence under this mechanism, the initial impact on reporting behavior will decline back to baseline, and this decline will be more rapid for income sources that have a lower autocorrelation. This is consistent with our findings, as seen in figure 4.

VII. Conclusion

This paper investigated the dynamic effects of audits on income reported in subsequent tax returns. Understanding these effects is important both from the perspective of quantifying the returns to the tax authority from an audit, and for assessing the mechanisms by which audits might influence taxpayer behavior. To answer this question, we exploited a random audit program run by the UK tax authority (HMRC) under which an average of around 4,900 individuals are randomly selected for audit each year. We used data on audits over the period 1998/1999–2008/2009, and we tracked responses on tax returns between 1998/1999 and 2011/2012.

We established three main results. First, we provided evidence of important *dynamic effects*, with the additional tax revenue over the five years postaudit equalling 1.5 times the direct revenue raised by audit. Second, we documented that a return to misreporting occurred more rapidly after audit for income sources that were less autocorrelated. Third, we showed that only those who were found to have made mistakes responded to the audit. Extending the standard model of rational tax evasion, we demonstrated that the observed dynamics are consistent only with audits revealing information to the tax authority, which makes misreporting certain income sources easier to detect for a period after the audit.

Our results have three main policy implications. First, taking dynamic effects into account substantially increases the estimated revenue impact of audits. The direct effect of an audit is (on average) £830, whilst the cumulative dynamic effect over the subsequent five years is £1,230, 1.5 times the direct effect. This suggests that the optimal audit rate should be substantially increased relative to the situation in which there are no dynamic effects. A back-of-the-envelope calculation suggests that the cost of an audit to the tax authority is around £2,500, so that even random audits are close to

¹⁶Full details and formalisation are provided in online appendix B.

¹⁷A formalisation of the following results is provided in online appendix B.

breaking even. For targeted audits, including dynamic effects raises the average return from around £6,000 to £15,000.

Second, the variation in dynamic effects observed across different income components alters the way in which targeted audits should be targeted: audits should focus more on individuals reporting types of income with the largest overall effects, combining immediate and dynamic effects. For example, the peak annual impact on reported self-employment income for each self-employed individual is over £1,000—higher than other components. This suggests focusing more on individuals reporting self-employment income. Likewise, although the maximum annual impact on pension income is lower, it is persistent, so there may be more incentive to target individuals believed to be underreporting pension income. The precise design of any targeting strategy must of course take into account how taxpayers would respond to the strategy, but for the tax authority the first step in designing any targeting strategy must be to know where the revenue is.

Third, there are implications for setting optimal reauditing strategies. Impacts for reported self-employment income and dividend income die away after about four years, so it might make sense to revisit these individuals around this time. In contrast, the impact on reported pension income seems to persist for at least eight years, implying that there is less of a need to reaudit these individuals so soon. Again, the responses of taxpayers to changes in audit strategy must be considered.

Our findings also highlight the importance of further study of the indirect effect of tax-compliance audits. One natural direction for further work would be to understand how the dynamic effects vary in the context of targeted audits, which are focused on individuals deemed likely to be noncompliant. A second avenue for exploration is the spillover effect of audits: does auditing taxpayers change the behavior of other taxpayers with whom they interact (Boning et al., 2020)? A third question is the extent to which cheaper “threat letters” can be used to maintain consistently high levels of compliance over the long term in the absence of high audit probabilities. A better understanding of these effects is crucial in determining optimal audit policy.

Finally, our results speak to the wider use of audits for public policy, whether it be to reduce corruption, improve public service delivery, or ensure environmental standards are met. A key lesson is that audits change future behavior but *how* that behavior changes depends on the likelihood of being caught in the future. Unless there are ongoing incentives to improve compliance—such as increased audit risk, increased penalties, or easier verification of misreporting—changes in reporting may be short-lived. However, a key tradeoff in public policy contexts is that individuals may be able to discontinue activities that are subject to audit if the strictness of enforcement is too high. This limits the compliance improvements achieved (Tulli, 2019), and it may have additional welfare costs as some valuable activities become more expensive (Gerardino et al., 2020) or do not take place (Lichand, 2016).

Appendices

Appendix A Additional Tables and Figures

TABLE A1.—RANDOM AUDIT LAGS AND DURATIONS

	Mean	Std. dev.	Median	75th	90th
Lag to audit start (months)	8.9	4.0	9	11	14
Audit duration (months)	5.3	6.6	3	7	13
Total time to audit end (months)	14.3	7.3	13	17	23

Annual averages for tax years 1998/1999–2008/2009. Includes all individuals with a completed random audit.

Source: Authors' calculations based on HMRC administrative data sets.

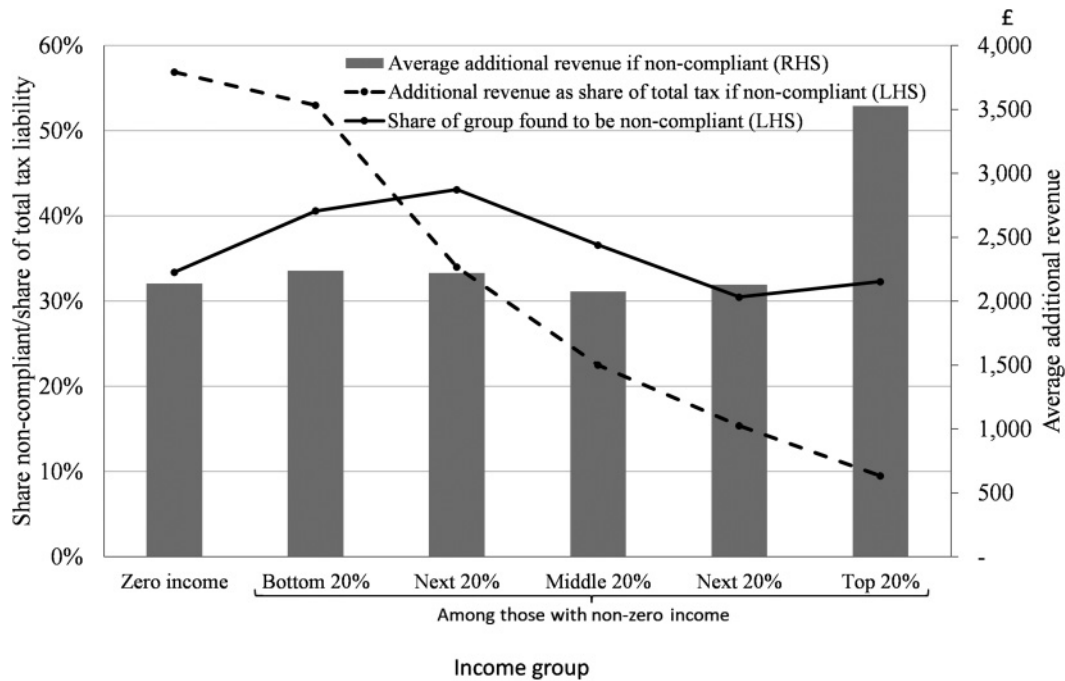
TABLE A2.—SAMPLE BALANCE (UNCONDITIONAL)

Years after audit		–5	–4	–3	–2	–1
Characteristics						
Female	Mean	.274	.276	.278	.282	.287
	Difference	–.006	–.004	–.002	–.001	–.002
	<i>p</i> -value	.221	.359	.606	.627	.863
Age	Mean	49.2	49.3	49.3	49.4	49.5
	Difference	.2	.3	.3	.2	.2
	<i>p</i> -value	.756	.586	.390	.610	.057
In London or SE	Mean	.333	.334	.335	.333	.331
	Difference	–.006	.001*	.004*	.002	.002
	<i>p</i> -value	.177	.025	.011	.281	.152
Has tax agent	Mean	.628	.614	.603	.589	.573
	Difference	.000	.002	.001	.003	.005
	<i>p</i> -value	.547	.508	.405	.396	.412
Survives	Mean	.624	.669	.728	.803	.892
	Difference	.032***	.039***	.047***	.050***	.050***
	<i>p</i> -value	.000	.000	.000	.000	.000
Income and tax totals						
Total taxable income	Mean	35,075	34,670	34,030	32,912	31,755
	Difference	881	492	403*	1,051*	1,095*
	<i>p</i> -value	.374	.157	.028	.012	.012
Total tax	Mean	9,646	9,539	9,321	8,979	8,635
	Difference	260	63	82	310	337*
	<i>p</i> -value	.539	.303	.055	.064	.027
Income components						
Employment	Mean	22,508	22,534	22,266	21,708	21,145
	Difference	–31	–136	180*	909**	721*
	<i>p</i> -value	.162	.371	.028	.006	.027
Self-employment	Mean	6,546	6,379	6,200	5,950	5,581
	Difference	356	328	173	99	200*
	<i>p</i> -value	.151	.174	.311	.106	.025
Interest and dividends	Mean	4,007	3,905	3,895	3,759	3,645
	Difference	–36	208	18	63	112
	<i>p</i> -value	.767	.432	.700	.578	.580
Pensions	Mean	3,493	3,542	3,561	3,562	3,531
	Difference	176	168	128	148	159
	<i>p</i> -value	.425	.478	.642	.307	.327
Property	Mean	869	844	811	769	726
	Difference	18	–2	37	47	31
	<i>p</i> -value	.813	.952	.576	.498	.134

“Years after audit” measures time relative to audit, or placebo audit for controls. “Mean” is the mean outcome in the control (not selected for audit) group across all years. “Difference” is the coefficient on the treatment dummy in a regression of the outcome on a treatment dummy. Treatment dummy equals 1 if taxpayer was selected by HMRC for a random audit. *p*-values are derived from an F-test that coefficients on interactions between treatment and tax year dummies are all zero in a regression of the outcome of interest on tax year dummies and interactions between treatment and tax year dummies. This is a stronger test than just testing the coefficient on treatment not interacted. “Survives” indicates presence in the data. Tests for all outcomes other than “survives” are conditional on survives = 1. Monetary values are in 2012 prices. Standard errors are clustered by taxpayer. **p* < .05, ***p* < .01, and ****p* < .001.

Source: Authors' calculations based on HMRC administrative datasets.

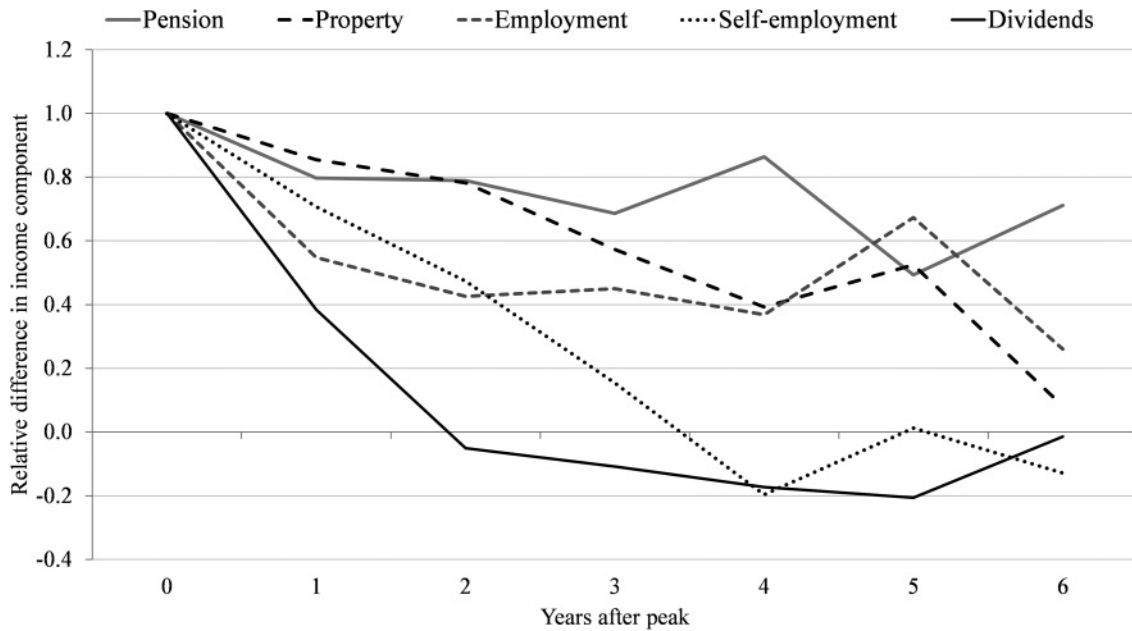
FIGURE A1.—NONCOMPLIANCE OVER THE PRIOR YEAR'S REPORTED INCOME DISTRIBUTION



Constructed using data on individuals who received a random audit of their self-assessment tax return for a tax year between 1998/1999 and 2008/2009. Income grouping is done based on previous year's reported income. A total of 16.2% of individuals report having zero income in the previous year. The remaining individuals are divided into five equal-sized bins based on their previous income: quintiles conditional on reporting nonzero income. "Share of group found to be non-compliant" is the share of individual taxpayers who are found to owe additional tax when audited. "Average additional revenue if non-compliant" is the average total tax in 2012 that was not reported among those individuals for whom some tax was not reported (the non-compliant). "Additional revenue as a share of total tax if non-compliant" is the additional tax owed divided by total tax owed, averaged across individual taxpayers who were non-compliant.

Source: Advani (2022).

FIGURE A2.—RELATIVE DYNAMICS BY INCOME SOURCE, AFTER REWEIGHTING BY CHARACTERISTICS



Sample includes individuals selected for a random audit between 1998/1999 and 2008/2009, and control individuals who could have been selected in the same years but were not. It uses tax returns from 1998/1999 to 2011/2012. Each line plots the point estimate for the difference in the average of a particular component of income between individuals who were and weren't audited, for different numbers of years after the peak impact for that income source, after reweighting individuals so that the distribution of observed characteristics matches that seen among the self-employed. This comes from dividing individuals into groups by sex, age band, and quartile of filing history. Observations are reweighted so that the distribution across these discrete cells is the same as for the self-employed. Point estimates for the treatment effect come from a weighted regression of each income component on dummies for years since audit (or placebo audit for controls), dummies for years since audit (or placebo audit for controls) interacted with treatment status, tax year dummies, and dummies for whether the taxpayer filed a return in each of the four years before audit, with audit status instrumented by selection for audit.

Source: Calculations based on HMRC administrative data sets.

REFERENCES

- Advani, Arun, "Who Does and Doesn't Pay Taxes?" *Fiscal Studies* 43:1 (2022), 5–22. 10.1111/1475-5890.12257
- Allingham, Michael, and Agnar Sandmo, "Income Tax Evasion: A Theoretical Analysis," *Journal of Public Economics* 1 (1972), 323–338. 10.1016/0047-2727(72)90010-2
- Alm, James, "What Motivates Tax Compliance?" *Journal of Economic Surveys* 33:2 (2019), 353–388. 10.1111/joes.12272
- Almunia, Miguel, and David Lopez-Rodriguez, "Under the Radar: The Effects of Monitoring Firms on Tax Compliance," *American Economic Journal: Economic Policy* 10:1 (2018), 1–38. 10.1257/pol.20160229
- Asatryan, Zareh, and Andreas Peichl, "Responses of Firms to Tax, Administrative and Accounting Rules: Evidence from Armenia," CESifo working paper 6754 (2017).
- Avis, Eric, Claudio Ferraz, and Frederico Finan, "Do Government Audits Reduce Corruption? Estimating the Impacts of Exposing Corrupt Politicians," *Journal of Political Economy* 126:5 (2018), 1912–1964. 10.1086/699209
- Becker, Gary, "Crime and Punishment: An Economic Approach," *Journal of Political Economy* 169 (1968), 176–177.
- Beer, Sebastian, Matthias Kasper, Erich Kirchler, and Brian Erard, "Do Audits Deter or Provoke Future Tax Noncompliance? Evidence on Self-employed Taxpayers," *CESifo Economic Studies* 66:3 (2020), 248–264. 10.1093/cesifo/izf018
- Bergolo, Marcelo, Rodrigo Ceni, Guillermo Cruces, Matias Giacobasso, and Ricardo Perez-Truglia, "Tax Audits as Scarecrows: Evidence from a Large-Scale Field Experiment," NBER working paper 23631 (2020).
- Bloomquist, Kim, "Incorporating Indirect Effects in Audit Case Selection: An Agent-Based Approach," *IRS Research Bulletin* (2013), 103–116.
- Blumenthal, Marsha, Charles Christian, and Joel Slemrod, "Do Normative Appeals Affect Tax Compliance? Evidence from a Controlled Experiment in Minnesota," *National Tax Journal* 54:1 (2001), 125–138. 10.17310/ntj.2001.1.06
- Bobonis, Gustavo J., Luis R. Cámara Fuertes, and Rainer Schwabe, "Monitoring Corruptible Politicians," *American Economic Review* 106:8 (2016), 2371–2405. 10.1257/aer.20130874
- Boning, William C., John Guyton, Ronald H. Hodge, Joel Slemrod, and Ugo Troiano, "Heard it through the Grapevine: Direct and Network Effects of a Tax Enforcement Field Experiment," *Journal of Public Economics* 190 (2020), 104261. 10.1016/j.jpubeco.2020.104261
- Burgherr, David, "The Costs of Administering a Wealth Tax," *Fiscal Studies*, 42:3–4 (2021), 677–697. 10.1111/1475-5890.12276
- Carrillo, Paul, Dina Pomeranz, and Monica Singhal, "Dodging the Taxman: Firm Misreporting and Limits to Tax Enforcement," *American Economic Journal: Applied Economics* 9:2 (2017), 144–164. 10.1257/app.20140495
- Choo, Lawrence, Miguel Fonseca, and Gareth Myles, "Lab Experiment to Investigate Tax Compliance: Audit Strategies and Messaging," HM Revenue and Customs Research Report 308 (2013).
- Cowell, Frank A., *Cheating the Government: The Economics of Evasion* (Cambridge, MA: MIT Press, 1990).
- DeBacker, Jason, Bradley Heim, Anh Tran, and Alexander Yuskavage, "Once Bitten, Twice Shy? The Lasting Impact of IRS Audits on Individual Tax Reporting," *Journal of Law and Economics* 61 (2018), 1–35. 10.1086/697683
- Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan, "Truth-Telling by Third-Party Auditors and the Response of Polluting Firms: Experimental Evidence from India," *Quarterly Journal of Economics* 128:4 (2013), 1499–1545. 10.1093/qje/qjt024
- , "The Value of Regulatory Discretion: Estimates from Environmental Inspections in India," *Econometrica* 86:6 (2018), 2123–2160.
- Dwenger, Nadja, Henrik J. Kleven, Imran Rasul, and Johannes Rincke, "Extrinsic and Intrinsic Motivations for Tax Compliance: Evidence from a Field Experiment in Germany," *American Economic Journal: Economic Policy* 8:3 (2016), 203–232. 10.1257/pol.20150083
- Erard, Brian, "The Influence of Tax Audits on Reporting Behaviour" (pp. 95–114), in Joel Slemrod (ed.), *Why People Pay Taxes: Tax Compliance and Enforcement* (Ann Arbor, MA: University of Michigan Press, 1992).
- Fellner, Gerlinde, Rupert Sausgruber, and Christian Traxler, "Testing Enforcement Strategies in the Field: Threat, Moral Appeal and Social Information," *Journal of the European Economic Association* 11:3 (2013), 634–660. 10.1111/jeea.12013
- Gemmell, Norman, and Marisa Ratto, "Behavioral Responses to Taxpayer Audits: Evidence from Random Taxpayer Inquiries," *National Tax Journal* 65 (2012), 33–58. 10.17310/ntj.2012.1.02
- Gerardino, Maria Paula, Stephan Litschig, and Dina Pomeranz, "Distortion by Audit: Evidence from Public Procurement," NBER working paper 23978 (2020).
- Kastlunger, Barbara, Erich Kirchler, Luigi Mittone, and Julia Pitters, "Sequences of Audits, Tax Compliance, and Taxpaying Strategies," *Journal of Economic Psychology* 30 (2009), 405–418. 10.1016/j.joep.2008.10.004
- Kleven, Henrik J., Martin Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez, "Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark," *Econometrica* 79 (2011), 651–692. 10.3982/ECTA9113
- Kleven, Henrik J., Claus Thustrup Kreiner, and Emmanuel Saez, "Why Can Modern Governments Tax So Much? An Agency Model of Firms as Fiscal Intermediaries," *Economica* 83:330 (2016), 219–246.
- Kolm, Serge-Christophe, "A Note on Optimum Tax Evasion," *Journal of Public Economics* 2:3 (1973), 265–270. 10.1016/0047-2727(73)90018-2
- Lichand, Guilherme, "Is Corruption Good for Your Health?" PhD thesis, Harvard University (2016). <https://dash.harvard.edu/bitstream/handle/1/33493343/LICHAND-DISSERTATION-2016.pdf>
- Long, Susan, and Richard Schwartz, "The Impact of IRS Audits on Taxpayer Compliance: A Field Experiment in Specific Deterrence," Paper presented at the Annual Meeting of the Law and Society Association, Washington DC (1987).
- Løyland, Knut, Oddbjørn Raaum, Gaute Torsvik, and Arnstein Øvrum, "Compliance Effects of Risk-Based Tax Audits," CESifo working paper 7616 (2019).
- Maciejovsky, Boris, Erich Kirchler, and Herbert Schwarzenberger, "Misperception of Chance and Loss Repair: On the Dynamics of Tax Compliance," *Journal of Economic Psychology* 28 (2007), 678–691. 10.1016/j.joep.2007.02.002
- Mascagni, Giulia, "From the Lab to the Field: A Review of Tax Experiments," *Journal of Economic Surveys* 32:2 (2018), 273–301. 10.1111/joes.12201
- Mittone, Luigi, "Dynamic Behaviour in Tax Evasion: An Experimental Approach," *Journal of Socio-Economics* 35 (2006), 813–835. 10.1016/j.socec.2005.11.065
- Naritomi, Joana, "Consumers as Tax Auditors," *American Economic Review* 109:9 (2019), 3031–3072. 10.1257/aer.20160658
- Perez-Truglia, Ricardo, and Ugo Troiano, "Shaming Tax Delinquents," *Journal of Public Economics* 167 (2018), 120–137. 10.1016/j.jpubeco.2018.09.008
- Pomeranz, Dina, "No Taxation Without Information: Deterrence and Self-Enforcement in the Value Added Tax," *American Economic Review* 105:8 (2015), 2539–2569. 10.1257/aer.20130393
- Pomeranz, Dina, and José Vila-Belda, "Taking State-Capacity Research to the Field: Insights from Collaborations with Tax Authorities," *Annual Review of Economics* 11:1 (2019), 755–781. 10.1146/annurev-economics-080218-030312
- Rosenbaum, Paul, and Donald Rubin, "Constructing a Control Group Using Multivariate Matched Sampling Methods that Incorporate the Propensity Score," *American Statistician* 39 (1985), 33–38.
- Rubin, Donald, "Using Propensity Scores to Help Design Observational Studies: Application to the Tobacco Litigation," *Health Services and Outcomes Research Methodology* 2 (2001), 169–188. 10.1023/A:1020363010465
- Sarin, Natasha, and Lawrence H. Summers, "Understanding the Revenue Potential of Tax Compliance Investment," NBER working paper 27571 (2020).
- Slemrod, Joel, "Tax Compliance and Enforcement," *Journal of Economic Literature* 57:4 (2019), 904–954. 10.1257/jel.20181437
- Slemrod, Joel, Marsha Blumenthal, and Charles Christian, "Taxpayer Response to an Increased Probability of Audit: Evidence from a Controlled Experiment in Minnesota," *Journal of Public Economics* 79:3 (2001), 455–483. 10.1016/S0047-2727(99)00107-3
- Slemrod, Joel, Brett Collins, Jeffrey L. Hoopes, Daniel Reck, and Michael Sebastiani, "Does Credit-Card Information Reporting Improve

- Small-Business Tax Compliance?" *Journal of Public Economics* 149 (2017), 1–19. 10.1016/j.jpubeco.2017.02.010
- Slemrod, Joel, and Shlomo Yitzhaki, "The Optimal Size of a Tax Collection Agency," *The Scandinavian Journal of Economics* 89:2 (1987), 183–192. 10.2307/3440063
- "Tax Avoidance, Evasion, and Administration," in A. J. Auerbach and M. Feldstein (ed.), *Handbook of Public Economics III* (Amsterdam: Elsevier, 2002).
- Tauchen, Helen, Ann Witte, and Kurt Beron, "Tax Compliance: An Investigation Using Individual Taxpayer Compliance Measurement Program (TCMP) data," *Journal of Quantitative Criminology* 9 (1993), 177–202. 10.1007/BF01071167
- Tulli, Andrea, "Sweeping the Dirt Under the Rug: Measuring Spillovers from an Anti-Corruption Measure," Job Market Paper (2019).
- Yitzhaki, Shlomo, "On the Excess Burden of Tax Evasion," *Public Finance Quarterly* 15 (1987), 123–137. 10.1177/109114218701500201
- Zamboni, Yves, and Stephan Litschig, "Audit Risk and Rent Extraction: Evidence from a Randomized Evaluation in Brazil," *Journal of Development Economics* 134 (2018), 133–149. 10.1016/j.jdeveco.2018.03.008