Tax Administration Research Centre

How long-lasting are the effects of audits?

Arun Advani
Institute for Fiscal Studies

William Elming
Institute for Fiscal Studies

Jonathan Shaw
Institute for Fiscal Studies

Discussion Paper: 011-15

How long-lasting are the effects of audits?

Preliminary. Please do not cite without authors′ permission

Arun Advani* William Elming† Jonathan Shaw*

March 5, 2015

Abstract

Understanding tax non-compliance and the effectiveness of strategies to tackle it is crucial for a modern tax authority. In this paper we study the indirect benefits of conducting audits, focusing on how the reported tax liability of audited individuals changes over time after an audit. We exploit data from a random audit program covering income tax self-assessment returns in the UK. We find that audits have a large and persistent impact on reported tax liability that reaches around 26 per cent on average by the fourth year following the tax year to which the audit relates.

*The Institute for Fiscal Studies (IFS), the Tax Administration Research Center (TARC) and University College London.

†IFS and TARC.

1 Introduction

The ongoing drive for a more efficient public sector increases the importance of understanding what determines tax non-compliance—and of having effective strategies to identify and recover unpaid revenue. Taxpayer audits are one widely-used example of such a strategy. Audits have a direct benefit in terms of additional revenue raised. There is, however, also potential for indirect benefits of audits.

Indirect effects take two forms: dynamic effects and spillover effects. Dynamic effects are changes in the future behaviour of the audited taxpayer. Spillover effects are changes in the behaviour of other taxpayers who know the audited taxpayer. These effects comes from a combination of updated information about the probability of an audit, updated information about the effectiveness of an audit, and updated information about the cost of an audit.

In this paper we focus on the dynamic effects of an audit. We study the tax returns of individuals, and examine how they change over time after an audit. Historically, tax authorities such as the Internal Revenue Service (IRS) have primarily focused on direct revenue maximization in selecting tax returns for their examination program (see Bloomquist (2013)). However, since most taxpayers pay taxes for many years, it is potentially important to understand the long term effects these audits have. This paper therefore seeks to quantify the amount of additional revenue received from an individual taxpayer in the years after audit. Understanding this is crucial in determining the total return from an audit, and hence in determining the optimal extent of enforcement.

We exploit a random audit program run by the UK tax authority, Her Majesty's Revenue and Customs (HMRC). The program we study is used by HMRC to estimate a model of non-compliance on income tax self-assessment. An average of 2,827 individuals are selected for a random audit each year, corresponding to a random audit probability of .03 per cent.[1] Taxpayers are selected for audit with equal probability among the population of self-assessment taxpayers.[2] After individuals are assigned to audit a few cases are deselected.[3] Thus, we cannot just compare audited individuals to non-audited individuals to isolate the causal effect of audits.

To find a suitable control group we exploit the fact that the rules for inclusion in the random audit sample remained broadly unchanged across the years we study. We therefore use individuals who were audited under the same program but further in the future as our control group. For example, we can compare the reported tax in 2006 for an individual audited in 2005 (i.e. the 'treatment

[1] In addition to the random audits, an average of 162,500 cases are selected for targeted audits per year. This corresponds to a targeted audit probability of 1.8 per cent.

[2] Individuals have to submit a self-assessment if they have a income which is not correctly taxed by withholding, e.g. if they are high earners; have capital gains; have income from UK land or property; are self-employed; are partners in a business partnership; or are company directors. See https://www.gov.uk/self-assessment-tax-returns/who-must-send-a-tax-return for more details about the self-assessment taxpayers eligibility criteria.

[3] This will happen if, for example, the tax authorities can verify all reported income using third party information.

effect' one year after the audit) to the reported tax in 2006 for an individual audited in 2007. As this 'control individual' had not yet been subject to an audit in 2006, comparing the two will isolate the causal effect of the audit one year after the audit.[4]

We focus on the population of randomly selected taxpayers, rather than using the targeted audits. This allows us to recover externally valid estimates of the average effect of an audit across the entire population of self-assessment taxpayers.

We have data on audits for 14 years between 1996/97 and 2009/10 and income tax returns for 16 years between 1996/97 and 2011/12. We currently use years 2004/05 to 2009/10: some additional information is required before we can use earlier years. This means we can follow up individuals for a substantial amount of time following the audits.

Our main results are summarised as follows. After the audit has taken place, there is a large and persistent impact on reported tax liability that reaches around 26 per cent on average by the fourth year following the tax year to which the audit relates. There are no substantial differences between treatment and control groups before the audit has taken place, suggesting that we have a valid control group.

Looking across different quantiles, we find that the effect of audits is felt across the whole distribution. We find a larger and more volatile effect at the .25 quantile. However, this may simply reflect the small base level of reported tax liability for this group.

The next section presents some of the related literature. Section 3 outlines the policy context, while section 4 describes the data and presents some summery statistics. Section 5 presents the method used, and discus the way in which we construct our control group. The results are presented in section 6. Section 7 concludes.

2 Related literature

Following the seminal paper by Allingham and Sandmo (1972) there has been a large literature that studies tax compliance, evasion and enforcement regimes (see Andreoni et al. (1998), Slemrod and Yitzhaki (2002) and Slemrod (2007) for more exhaustive reviews of the literature). In general, the literature can be divided into three subgroups: a literature that uses individual taxpayer (micro) data, a literature that relies on aggregate data, and finally a literature that uses laboratory experiments.

2.1 Micro data

Within the literature using micro data, a number of previous papers have investigated the deterrence effect of audits. Kleven et al. (2011) analyze the effect of

[4]This–of course–assumes that there are no differences between the two individuals other than the timing of the audit. We return to this important point below.

a tax enforcement field experiment in Denmark. Half of a sample of income tax filers were randomly selected for audit, while the other half were not audited. The following year, letters threatening an audit were randomly assigned to tax filers in both groups. Declarations of self-reported and third-party reported income were then followed up. They find that prior audits and threat-of-audit letters have significant effects on self-reported income, but no effect on third-party reported income.

Gemmell and Ratto (2012) investigate behavioural responses to taxpayer audits using earlier versions of the random audit data we use. They distinguish between taxpayers found to be non-compliant and those found to be compliant, arguing that the former are likely to increase their subsequent compliance while the latter could reduce their compliance.[5] However, this distinction between compliant and non-compliant taxpayers is endogenous, making it hard to interpret the comparison with an unconditionally randomly selected control group as causal.

There are few other comparable studies of the behavioural effects of audits on future income reporting. Long and Schwartz (1987) find that audits in 1969 had no effect on average non-compliance in 1971 and only a small effect on the frequency of non-compliance. Erard (1992) uses IRS data to assess the effect of audits on taxpayers' compliance in the following year. His findings indicate that evaders increase the tax declared following an audit, though the results are not conclusive. The effects he finds cannot be solely attributed to the effect of the audit, and the results are sensitive to the way selection issues are dealt with in the estimation.

Tauchen et al. (1993) estimate the general deterrence effect of audits on the amount of income that taxpayers choose to report on their tax returns. They use a stratified random sample of approximately 50,000 individual tax returns (from the 1979 Taxpayer Compliance Measurement Program) and combine the returns with IRS administrative records for District Offices and 1980 Census data at the five-digit zip code level. They find weak evidence that higher audit rates are associated with greater levels of compliance. They find that a 10 per cent increase in audits lead to a 2.3 per cent increase in reported income for their high-income group, with no significant effects on the other income groups. Further they find that the general deterrence effects of audits are over $2 for every $1 of direct revenue yield.

Bergman and Nevarez (2006) use VAT tax return information and enforcement data to determine the effect of audits on subsequent compliance of taxpayers in Argentina and Chile. They find no evidence that audits increase individual compliance.

While the papers above focus on the effect of actual audits on reporting behaviour, Slemrod et al. (2001) and Agostini and Martnez (2014) studies the effect of tax authority letters containing *threat* of audits. Slemrod et al. (2001) compare the change in reported tax for 1,724 randomly selected Minnesota tax-

[5]This could be the case if they are audited and found compliant, but they were in fact non-compliant.

payers, who received a letter with a threat of audit, relative to a control group that did not receive such a letter. They find that low and middle-income taxpayers in the treatment group on average increased tax payments compared to the previous year. The effect is only significant for those with more opportunities to evade taxes (i.e. those with self-employment or farm income). For the high-income treatment group, however, the conclusions are radically different; upon receiving the threat-of-audit letter they decrease their reported tax liability relative to the control group. Slemrod et al. (2001) argue that this is because higher income taxpayers see the audit threat as the beginning of a negotiation over actual tax liability.

Agostini and Martnez (2014) study the impact of a tax enforcement programme implemented by the Chilean Internal Revenue Service, where letters requiring information about diesel purchases and use and vehicle ownership were sent to around 200 firms in 2003. They find that firms receiving a letter decreased their diesel tax credits by around 10 per cent.

2.2 Aggregate data

In addition to the literature that uses micro data, there is a literature that relies on aggregate data to examine the effects of audits. Dubin et al. (1990) uses the Annual Report of the Commissioner of Internal Revenue and the Statistics of Income for the years 1977-1986 to create a state-level data set. They set up three different linear equations where 'reported tax per return', 'returns filed per capita' and 'assessed tax liability per return' are the dependent variables. As explanatory variables they use the state income tax, the audit rate, per-capita income and other socioeconomic variables. They find that the audit rate had a significant and positive effect on reported tax and assessed liabilities per return. Hence they conclude that audits are an effective deterrent against taxpayer noncompliance. They then use the estimated linear relationships to calculate the counterfactual scenario where audit rates are kept constant at their 1977 level throughout the period. They find that had the audit rate remained constant at the 1977 level throughout the period from 1977-1986 then the total reported taxes would have been greater by 15.6 billion dollars in 1986 (corresponding to roughly 4 per cent of total reported tax).

Plumley (1996) uses a ten-year (1982-91) panel data set aggregated to the state level. He finds that the general deterrence effect of audits for the general population is about 11 times as large as the adjustments proposed by the audits themselves. Ali et al. (2001) use data from the Annual Report of the Commissioner of Internal Revenue Service and the Data Book for 1980 to 1995 to investigate the relationship between taxpayer compliance, audit rates and penalties if detected. They find that both the audit and penalty rate are effective deterrents of non-compliance. Further, they find that the effectiveness of these deterrents is increasing in income. Overall, they find that compliance increases with income, though at a decreasing rate.

Dubin (2007) follows the method of Dubin et al. (1990) and uses state-level data covering the years 1988-2001 to empirically test whether measurable activ-

ities of the IRS Criminal Investigation Division (CI) affect taxpayer compliance. He finds that CI activities have a measurable and significant effect on voluntary compliance and that incarceration and probation (rather than fines) have the largest effect.

2.3 Experimental data

Finally, a number of papers use experimental set-ups to test how different circumstances surrounding audits affect reporting behaviour. The papers by Alm and Mckee (2004), Fortin et al. (2007) and Alm et al. (2009) investigate how interactions among taxpayers affects the deterrence effects of audits. Alm and Mckee (2004) studies compliance behaviour when returns are selected for audit based on the deviation of each individuals tax report from the average reported tax of other taxpayers in their cohort. Thus, the optimal strategy of the participants in the experiment would be to cooperate and reach the zero compliance equilibrium, where no one is audited and all have the maximum possible gain.[6] They find that participants struggle to reach this equilibrium, but that pre-game communication facilitates coordination.

In a similar spirit, Fortin et al. (2007) studies the impact of social interactions on tax evasion using experimental data. Their experimental results provide evidence of fairness effects. Specifically, for a given gross income and personal tax rate, the individual will report less if they feel they are being treated unfair by the tax system (e.g. if they pay a higher tax rate than other participants in their group). Alm et al. (2009) investigate the effects of information dissemination concerning enforcement and compliance behaviour of others on the tax reporting behaviour of individual taxpayers in laboratory experiments. They find that taxpayers will respond to wide ranging information sources that report the enforcement effort and estimate the general deterrence effect on compliance to be 4.4 times the direct effect.

While the experimental papers above study the interactive effects of audits the papers of Friedland et al. (1978), Kirchler et al. (2007) and Kastlunger et al. (2009) examine the intertemporal effects of audits. Friedland et al. (1978) investigate the tax evasion behaviour of 15 participants in a game-simulation context. They find that large fines tend to be more effective deterrents than frequent audits. Kirchler et al. (2007) investigate the effectiveness of audit probabilities and sanctions in a dynamic setting focusing on the time lag between audits. They find that compliance decreases immediately after an audit, suggestive of a 'bomb crater effect'.[7] Contrary to Friedland et al. (1978), Kirchler

[6] In some of their experiments the audit rule based on differences in the reported tax from the mean is augmented with a random audit rule such that there is a positive probability of audit regardless of the reporting strategy employed.

[7] The 'bomb crater effect' refers to the idea that individuals might perceive the risk of being audited to fall immediately after an audit. The name originates from preference of WW1 soldiers to hide out in bomb craters, believing that it was unlikely that a bomb would strike the exact same place again (see Mittone (2006)). A competing explanation for the decline in reported tax following an audit is the mechanism of loss repair: experiencing an audit may make taxpayers "want to evade more in the future in an attempt to 'get back' at

et al. (2007) find that the effect of sanctions on compliance to be relatively less important than higher audit probabilities. They do, however, find that larger sanctions are associated with the tendency of participants to repair their losses following and audit.

Kastlunger et al. (2009) use a laboratory experiment where participants file taxes 60 times (i.e. participants' tax lifecycle). They then subject participants to different patterns of audits in two different studies. The first study focuses on the immediate reaction in reported income following an audit. Similarly to Kirchler et al. (2007) they find a strong decrease in compliance in the period following the audit. Further they find that this is most likely to be the 'bomb crater effect' rather than loss-repair tendencies. The second study investigates the effect of the timing of the audits in participant's tax lifecycle. They find that early, rather than late, audit experiences in participants tax lifecycle lead to increased compliance.

Lastly, Choo et al. (2013) investigate the compliance behaviour of 92 self-assessed taxpayers. They find that increasing the audit rate had no significant effect on the number of people attempting to evade, and no significant effect on the amount of evasion among those who did evade. Further, the main finding of their experiment is that the compliance levels by self-assessed taxpayer participants are extremely high and non-responsive to changes in audit rates. Further, post-experiment questionnaire data shows compliant participants are driven by strong norms of honesty. In contrast, non-compliant participants are driven by profit maximisation. Their results suggest that the deterrence power of random audits is quite limited. Finally, they do not find any evidence that compliance levels drop immediately after an audit (the 'bomb crater effect'), neither for compliant or non-compliant participants. This is in contrast to the experimental evidence from Kirchler et al. (2007) and Kastlunger et al. (2009), which used student participants.

3 Policy context

Income tax is the largest of all UK taxes, contributing 26.2 per cent of total government receipts in 2012-13. Most (but not all) sources of income are subject to income tax, including earnings, retirement pensions, income from property, interest on deposits in bank accounts, dividends, and some 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.

A total of around 30 million individuals in the UK pay income tax each year. Out of these, around 8 million are subject to self-assessment and required to submit a tax return. These tend to be individuals with forms of income not subject to withholding or for whom the tax system struggles to calculate and

the tax agency" (p. 844 in Andreoni et al. (1998)).

Table 1: Average number of cases selected for income tax audits

	Random	Random audit probability	Targeted	Targeted audit probability
Individuals	2,827	.0003	152,585	.0176
Partnerships	261		8,172	
Trusts	61		1,797	
Total	3,150		162,553	

Notes: Annual averages for tax years 1996/97 to 2009/10. Data for audits is not available after 2009/10. Source: calculations based on HMRC administrative datasets.

withhold the right amount of tax. It includes self-employed individuals, those with very high incomes, company directors, landlords and many pensioners.

Since incomes covered by self-assessment tend to be harder to verify, there is a significant risk of non-compliance. As a result, HMRC carries out audits each year to deter non-compliance and recover lost revenue. HMRC runs two types of income tax audits. Targeted audits are based on perceived risks of non-compliance. Random audits are used to ensure that all self-assessment taxpayers face a positive probability of being audited, as well as to collect information about the scale of non-compliance and predictors of non-compliance.

The timeline for the audit process is as follows. The tax year ends on 5 April. In April or May, HMRC issues a notice to file to taxpayers who they believe need to submit a tax return.[8] Cases to be subjected to a random audit are provisionally selected from the population of individuals issued with a notice to file. The main filing deadline for taxpayers is 31 January the following calendar year. This is the date by which individuals must submit their tax return. HMRC then deselects some cases from random enquiry.[9] At the same time, targeted audits are selected on the basis of the information provided in self-assessment returns and other intelligence. Individuals selected for a random audit are selected before decisions about targeted audits are made, and cannot be also selected for a targeted audit. Individuals selected for a random audit cannot also be selected for a targeted audit in the same tax year because a random audit is worked in exactly the same way as targeted full audits. Audits must be opened within a year of the date when the return was filed (until 2007/08, it was a year from the 31 January filing deadline for returns filed on time). Taxpayers subject to an audit are informed of this when the audit is opened but they are not told whether it is a random or targeted audit.

Table 1 shows the average number of cases HMRC have selected for income tax audits per year over the period 1996/97 to 2009/10. HMRC selected an average of 3,150 cases for random audits and 162,553 cases for targeted audits per year. The vast majority of these are audits of individuals (as apposed to partnerships and trusts). For individuals, the corresponding probabilities of

[8]This excludes most first-time filers, which HMRC doesn't yet know about.

[9]This deselection implies that we cannot simply compare 'randomly' audited individuals to non-audited individuals.

being selected for an audit are .03 per cent (three in 10,000) for random audits and 1.8 per cent for targeted audits.

Figure 1: Number of individuals selected for income tax audits on

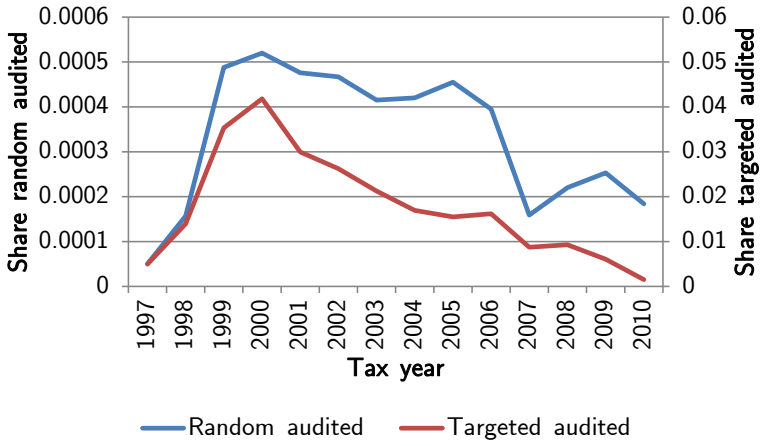

Source: calculations based on HMRC administrative datasets.

Figure 1 shows that there has been considerable variation in the audit probability for individuals over time. Audit rates peak in 1999/2000 at .05 per cent (five in 10,000) for random audits and 4.2 per cent for targeted audits, and then fall back substantially. The latest publicly-available estimates from the random audits programme described in HMRC (2013) suggest that 27 per cent of self-assessment taxpayers under-declared their tax liabilities in 2009/10.

4 Data

We exploit data on 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 numbers and tax year.

Audit records come from CQI (Compliance Quality Initiative), an operational HMRC dataset that records audits made into income tax self-assessment and corporation tax self-assessment returns. It does not include audits by the HMRC's Large Business Service, Special Investigations or Employer Compliance Reviews. It includes operational information about the audits, such as

Table 2: Lag in days between tax return filing and audit start

	Mean	Std. dev.
Random audits	239	92
Targeted audits	292	116

Notes: Annual averages for tax years 1996/97 to 2009/10. Includes all individuals under income tax self-assessment. Source: calculations based on HMRC administrative datasets.

start and end dates, and audit outcome (size of any correction, penalties and interest). There are also limited details about the taxpayer. There are around 50 variables in the data made available to us and, in the main table, there is one observation for each audit. We have CQI data covering audits for tax years 1996/97 to 2009/10.

We track individuals following the audit using information from tax returns. It is important to recognise that we have no way to identify actual compliance behaviour in the follow-up (i.e. subsequent to the initial audit). The number of random audit taxpayers that are re-audited is far too small for it to be possible just to focus on them.

Information from tax returns comes from two sources: SA302 and Valid View. The SA302 dataset contains information that is sent out to taxpayers summarising income and tax liability (SA302 forms). It is derived from self-assessment tax returns, which are put through a tax calculation process. There are a total of around 150 variables but the range of information it contains is relatively limited, e.g. there is no information about turnover and expenses for self-employed individuals, demographics or filing data. For each tax year, it contains one observation per self-assessment taxpayer. We currently have access to it for tax years 2004/05 to 2011/12.

SA302 is supplemented by variables drawn from Valid View, a dataset that provides information taken more directly from individuals' tax returns. Valid View includes variables relating to detailed income sources and tax liabilities, some demographics and filing information, but few income and tax totals of the form found in SA302. The dataset was assembled for internal HMRC analytical purposes and contains about 700 variables. For each year, there is one observation per taxpayer. We have access to it for tax years 1996/97 to 2011/12.

SA302 and, to a lesser extent, Valid View are both updated during the course of the tax year as new tax returns come in. The versions we are using correspond to the October following the end of the tax year to which the extract relates.

Table 2 shows the average lag in days between the date the tax return was filed and the date at which the audit was started: it is around eight months for random audits and almost 10 months for targeted audits. Figure 2 shows the distribution of random audit durations. The majority take less than six months, but there is a long tail of audits lasting a year or more. This substantial lag between when an audit is opened to when it is settled is important to keep in mind when interpreting the result later on.

Table 3 summarizes the outcomes of random audits on individuals. It shows

Figure 2: Distribution of audit durations

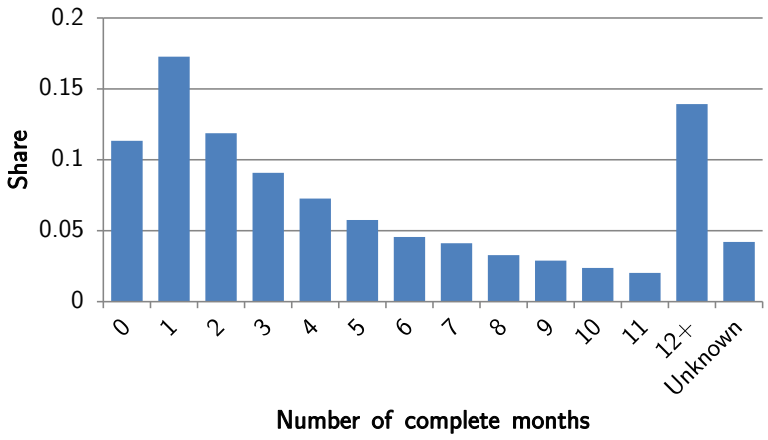

Source: calculations based on HMRC administrative datasets.

Table 3: Random audit outcomes

	Mean	Std. dev.
Correct	.529	.499
Incorrect (no underpayment)	.112	.316
Non-compliant (underpayment)	.359	.480
Has interest	.259	.438
Interest if positive (£)	442	1660
Has penalty	.064	.245
Penalty if positive (£)	1073	1979
Has additional tax	.363	.481
Additional tax if positive (£)	1949	8004
Has late penalty	.001	.034
Late penalty if positive (£)	883	2303
Return adjusted	.017	.128

Notes: Annual averages for tax years 1996/97 to 2009/10. Includes all individuals who faced a random audit. Source: calculations based on HMRC administrative datasets.

that 53 per cent of returns on average are found to be correct, 11 per cent are incorrect but involve no underpayment of tax and 36 per cent are non-compliant with a tax underpayment. In addition, 26 per cent have interest to pay and 6 per cent face a penalty. For cases where the return is incorrect, there are two possible types of settlement: a return adjustment (where the tax return is amended) and a contract settlement (with agreement between the parties as to how much additional tax should be paid but with no return adjustment). Less than 2 per cent of cases have return adjustment. This means that, in the vast majority of cases, the tax return information reflects the return as filed.

5 Method

Key to our approach is being able to identify a group of tax returns that are suitable to act as controls for audited cases. We focus on random rather than targeted audits because it enables us to say something about the population of self-assessment taxpayers rather than just the selected group that faces a targeted audit.

To isolate the dynamic effects of audits, we cannot simply compare individuals who experienced a random audit to individuals who were not subject to a random audit. As mentioned above, some individuals are deselected from the random audits program. If we were to compare the two groups we would not only isolate the effect of the 'random' audits, but also the effect of the deselection rules (as we do not know whether non-audited individuals would have been deselected, had they been assigned to receive a random audit).

To construct a suitable control group for the random audits we therefore use individuals who are randomly audited in the future. For example, controls for

returns audited in 2005 might be 2005 returns for taxpayers audited in 2006 (i.e tax reported prior to their audit). So long as the audit process is fairly stable over time, these future-audited observations should form a valid control group. We consider three control groups constructed in this way:

1. A rolling control group in which control observations for a given year are individuals who will be randomly audited the following year.

2. A maximal control group in which control observations for a given year are individuals who will be randomly audited at any point in the future.

3. A fixed control group in which control observations for a given year are individuals who will be randomly audited a fixed number of years in the future.

In the results below, we focus on the rolling control group.

One potential issue with using future-audited individuals as controls is survivor bias: for a given outcome horizon, individuals need to survive as self-assessment taxpayers different numbers of years to be in the treatment group compared to the control group. For example, for outcomes two years after the audit, treated individuals have to appear for three years (the audit year and two years later), while control individuals only have to appear one year before audit and in the audit year. This problem becomes worse the further into the future we consider. To control for this possible difference between our control and treatment group, we restrict both groups to have been in the sample equally long.

6 Results

The outcome of interest we focus on is the reported tax liability from the self-assessment return. For simplicity, we begin by presenting results for a single year. Table 4 shows the results of running Poisson regressions of reported tax liability on a treatment dummy and a constant for the year 2006. We do this separately for each horizon (two years before the audit, one year before, and so on). We use Poisson regression because tax liability is highly skewed so we want to fit a model of the form:

$$y_{it} = \exp(\alpha + \beta D_i + u_{it})$$

We prefer Poisson regression to using OLS on the log of tax liability as it allows for zeros in the data. We calculate robust standard errors to relax the usual Poisson assumption that the mean is equal to the variance.

In order to be confident that treatment and control groups are balanced, we hope to see that there is no effect of the audit before the audit has happened. This is exactly what we observe for the two years before the the tax year to which the audit relates: the there is basically no difference between treatment and control groups and the confidence intervals straddle zero.

Table 4: Impact of audit on reported tax liability: 2006 audits

Years since audit	-2	-1	0	1	2	3	4
Treated	.024	-.095	.125	.099	.210**	.255***	.215**
Std. Err.	(.106)	(.096)	(.075)	(.081)	(.071)	(.076)	(.076)
Constant	8.647***	8.659***	8.545***	8.545***	8.565***	8.606***	8.644***
Std. Err.	(.069)	(.078)	(.054)	(.054)	(.055)	(.057)	(.059)
N	2346	2631	4102	4094	3312	2929	2662

Notes: Based on Poisson regressions of reported tax liability on a treatment dummy and a constant. The regressions are run separately for different numbers of years since audit. * $p < .05$, ** $p < .01$, *** $p < .001$. Source: calculations based on HMRC administrative datasets.

Table 5: Impact of audit on reported tax liability: audit in any year

Years since audit	-2	-1	0	1	2	3	4
Treated	.059	.032	.01	.042	.131***	.166***	.233*
Std. Err.	(.038)	(.027)	(.022)	(.027)	(.034)	(.039)	(.102)
Constant	8.599***	8.621***	8.613***	8.631***	8.629***	8.665***	8.733***
Std. Err.	(.043)	(.036)	(.028)	(.03)	(.03)	(.031)	(.037)
N	7851	12559	19589	19836	18585	17286	16042

Notes: Based on Poisson regressions of reported tax liability on a treatment dummy, year dummies and a constant. The regressions are run separately for different numbers of years since audit. Standard errors are clustered by taxpayer. * $p < .05$, ** $p < .01$, *** $p < .001$. Source: calculations based on HMRC administrative datasets.

Figure 3: Impact of audit on reported tax liability: audit in any year

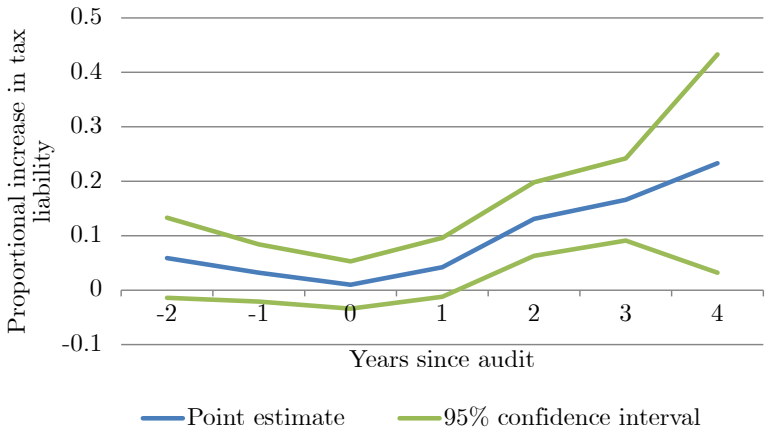

Source: calculations based on HMRC administrative datasets.

After the audit in 2006, the effect of the audit on reported tax liability climbs steadily, reaching a parameter value of around .22 after four years. Given the regression specification, this corresponds to an effect of 24 per cent ($e^{.215}-1$). It is not surprising that the effect of the audit increases over time, at least initially. This is because it takes some time for the audits to be started, and when they do get started, it takes some time to complete them (see Table 2 and Figure 2 above). But the size of the effect is considerable and persists up to four years out.

Table 5 and Figure 3 show the same results but now estimated across all years of audit data. We include year dummies in the regressions to control for year effects (excluded from the results table). These dummies are de-meaned so the constant in the regression still has the interpretation of average reported tax due in the control group across all years. We cluster standard errors to allow for the fact that different years' returns for the same individual may appear in the same regression. The patterns are consistent: no effect before the audit and a steadily climbing effect after the audit, reaching 26 per cent ($e^{.233}-1$) after four years. It should, however, be emphasised that the effect in the fourth year following the audit are estimated with considerable uncertainty as shown by the divergent confidence bands in Figure 3.

Figure 4: Impact of audit on quantiles of reported tax liability

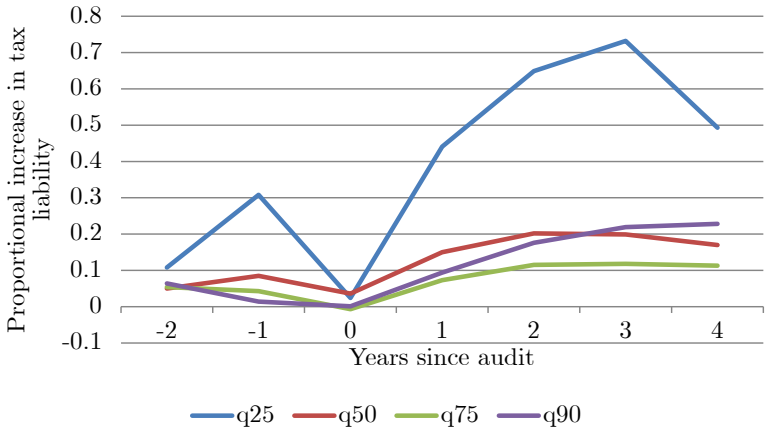

Source: calculations based on HMRC administrative datasets.

In Table 6 and Figure 4 we show what happens at different quantiles of the

Table 6: Impact of audit on quantiles of reported tax liability

	Quantile	-2	-1	0	1	2	3	4
Treated	q25	.108	.308**	.024	.441***	.649***	.732***	.493***
Std. Err.	q25	(.17)	(.118)	(.126)	(.123)	(.094)	(.106)	(.099)
Treated	q50	.05	.085	.036	.15***	.202***	.199***	.17***
Std. Err.	q50	(.051)	(.056)	(.04)	(.033)	(.037)	(.033)	(.028)
Treated	q75	.054	.043	-.007	.073**	.115***	.118***	.113***
Std. Err.	q75	(.041)	(.035)	(.021)	(.023)	(.023)	(.03)	(.028)
Treated	q90	.064	.014	.001	.094*	.176***	.219***	.228***
Std. Err.	q90	(.062)	(.059)	(.041)	(.045)	(.032)	(.045)	(.051)
N		7851	12559	19589	19836	18585	17286	16042

Notes: Based on quantile regressions of log tax liability on a treatment dummy, year dummies, and a constant, run separately for different numbers of years since an audit and different quantiles. * $p < .05$, ** $p < .01$, *** $p < .001$. Source: calculations based on HMRC administrative datasets.

distribution. These results are from quantile regressions of log tax liability on a treatment dummy, year dummies, and a constant, run separately for different numbers of years since an audit and different quantiles. These results show that the effect of audits is felt fairly evenly across the distribution. The .25 quantile is considerably more volatile than the other quantiles, a consequence of the low base.

7 Conclusion

In this paper we have investigated how long-lasting is the effect of audits on reported tax liability. This is an important question from the perspective of quantifying the returns to an audit, as well as understanding the mechanisms by which audits might influence taxpayer behaviour.

To answer this question we exploit a random audit program run by the UK tax authority (HMRC), where an average of 3,150 cases are selected for random audits each year. We use data on audits and individual income tax returns for the tax years 2004/05 to 2009/10.

Our results suggest that there is a large and persistent impact of audits on reported tax liability that reaches around 26 per cent by the fourth year following the tax year to which the audit relates. This emphasises the importance of taking the indirect revenue effects into account when deciding on the optimal enforcement strategy.

Future research should focus on quantifying all indirect benefit and comparing them coherently to the direct benefits from an audit, and the costs of conducting an audit.

References

Agostini, C. A. and C. Martnez, "Response of Tax Credit Claims to Tax Enforcement: Evidence from a Quasi-Experiment in Chile," *Fiscal Studies*, 2014, *35* (1), 4165.

Ali, M. M., H. W. Cecil, and J. A. Knoblett, "The effects of tax rates and enforcement policies on taxpayer compliance: A study of self-employed taxpayers," *Atlantic Economic Journal*, June 2001, *29* (2).

Allingham, M. G. and A. Sandmo, "Income tax evasion: A theoretical analysis," *Journal of Public Economics*, 1972, *1*, 323–338.

Alm, J. and M. Mckee, "Tax compliance as a coordination game," *Journal of Economic Behavior & Organization*, 2004, *54*, 297312.

_ , B. R. Jackson, and M. Mckee, "Getting the word out: Enforcement information dissemination and compliance behaviour," *Journal of public economics*, 2009, *93*, 392402.

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

Bergman, M. and A. Nevarez, "Do Audits Enhance Compliance? An Empirical Assessment of VAT Enforcement," *National Tax Journal*, 2006.

Bloomquist, K. M., "Incorporating Indirect Effects in Audit Case Selection: An Agent-Based Approach," *Forthcoming: IRS Research Bulletin*, 2013, *1500.*

Choo, L., M. A. Fonseca, and G. Myles, "Lab Experiment to Investigate Tax Compliance: Audit Strategies and Messaging," *HM Revenue and Customs Research*, 2013, *Report 308.*

Dubin, J. A., "Criminal Investigation Enforcement Activities and Taxpayer Noncompliance," *Public Finance Review*, 2007, *35* (3), 500–529.

_ , M. J. Greatz, and L. L. Wilde, "The effect of audit rates on the federal individual income tax, 1977-1986," *National Tax Journal*, December 1990, *43* (4), 395–409.

Erard, B, "The influence of Tax audits on reporting behaviour," *Why People Pay Taxes: Tax Compliance and Enforcement*, 1992.

Fortin, B., G. Lacroix, and M. Villeval, "Tax evasion and social interactions," *Journal of Public Economics*, 2007, *91*, 20892112.

Friedland, N., S. Maital, and A. Rutenberg, "A simulation study of income tax evasion," *Journal of Public Economics*, 1978, *10*, 107–116.

Gemmell, N and M Ratto, "Behavioral responses to taxpayer audits: Evidence from random taxpayer enquiries," *National Tax Journal*, 2012.

HMRC, "Measuring tax gaps: 2013 edition," 2013.

Kastlunger, B., E. Kirchler, L. Mittone, and J. Pitters, "Sequences of audits, tax compliance, and taxpaying strategies," *Journal of Economic Psychology*, 2009, *30*, 405418.

Kirchler, E., B. Macijovsky, and H Schwarzenberger, "Misperception of chance and loss repair: On the dynamics of tax compliance," *Journal of Economic Psychology*, 2007, *28*, 678691.

Kleven, H. J., M. B. Knudsen, C. T. Kreiner, S. Pedersen, and E. Saez, "Unwilling or unable to cheat? Evidence from a tax audit experiment in Denmark," *Econometrica*, 2011.

Long, S and R. Schwartz, "The Impact of IRS Audits on Taxpayer Compliance: A Field Experiment in Specific Deterrence," *Annual Law and Society Association Meeting 1987*, 1987.

Mittone, L., "Dynamic behaviour in tax evasion: An experimental approach," *The Journal of Socio-Economics*, October 2006, *35* (5), 813835.

Plumley, A. H., "The Determinants of Individual Income Tax Compliance Estimating the Impacts of Tax Policy, Enforcement, and IRS Responsiveness," *Internal Revenue Service Publication*, 1996, *1916*, Rev. 11–96.

Slemrod, J., "Cheating Ourselves: The Economics of Tax Evasion," *Journal of Economic Perspectives*, Winter 2007, *21* (1), 25–48.

_ **and S. Yitzhaki**, "Tax avoidance, evasion and administration," *Handbook of Public Economics*, 2002, *3.*

_ **, M. Blumenthal, and C. Christian**, "Taxpayer Response to an Increased Probability of Audit: Evidence from a Controlled Experiment in Minnesota," *Journal of Public Economics*, 2001.

Tauchen, H. V., A. D. Witte, and K. J. Beron, "Tax compliance: An investigation using individual taxpayer compliance measurement program (TCMP) data," *Journal of Quantitative Criminology*, 1993, *9* (2).