

Offshore Voluntary Disclosure Schemes: A Preliminary Analysis

*Matthew Gould and Matthew D. Rablen,
Brunel University London, United Kingdom*

1. Introduction

In recent years, tax authorities around the world have begun to enforce tax rules on offshore funds with forms of enforcement that combine aggressive information acquisition (but through non-audit means) with amnesty-type incentives for voluntary disclosure. This new form of enforcement may be broadly characterized by a two-stage process. In the first stage, the tax authority acquires (non-audit) information on the offshore assets of a set of taxpayers. In the second stage, the tax authority communicates the acquisition of information with a set of taxpayers (which will include, but may be larger than, the set of taxpayers on which it has information), and offers a one-off and time-limited opportunity to make a voluntary disclosure through a facility that offers overt incentives for honesty (in the form of lower fine or interest rates). We term schemes of this form Incentivized Offshore Voluntary Disclosure Schemes (IOVDS).

IOVDS mix together what have so far been three largely distinct strands of literature. The first is a literature that examines tax amnesties, which offer taxpayers reduced penalties if they wish to revise (upwards) their past tax returns (see, e.g., using theoretical (e.g., Andreoni, 1991; Malik and Schwab, 1991; Stella, 1991), empirical (e.g., Alm and Beck, 1993) and experimental (Alm, *et al.*, 1990) methods (Franzoni 2000)). IOVDS differ from an amnesty, however, in that the acquisition of information by the tax authority alters the taxpayer's beliefs over the likelihood of detection. Voluntary disclosure, therefore, takes place in the shadow of a credible threat of sanctions for those who choose to under-disclose (or make no disclosure at all). In contrast, an amnesty provides no new information to the taxpayer. Accordingly, a fully rational individual with perfect information—the type of individual considered in the standard economic model of tax compliance (e.g., Allingham and Sandmo, 1972)—might choose to make a positive disclosure in an IOVDS, whereas it is well known that such an individual would not participate in an amnesty (Andreoni, 1991; Malik and Schwab, 1991).¹

The second strand of literature is that on the optimal design of audit rules for the probability of audit conditional upon an observed income disclosure (e.g., Reinganum and Wilde, 1985; Mookherjee and Png, 1989; Sanchez and Sobel, 1993) and other signals correlated with income (Scotchmer, 1987; Macho-Stadler and Pérez-Castrillo, 2002). An IOVDS differs from a pure audit program in that, for such programs, the audit cost must be sunk before information leading to the recovery of unpaid tax can be found. In contrast, under an IOVDS, the tax authority may find itself knowing of the existence of unpaid tax without having carried out any audits. It may, nevertheless, still need to perform costly audits (purely for the purposes of formal legal verification) if it wishes to levy fines on the unpaid tax. Schemes that facilitate voluntary disclosure are of interest in these circumstances, for they provide the tax authority a means to avoid these additional verification costs.

The final strand of literature is that on the design of whistleblower incentive schemes, such as that run by the United States Internal Revenue Service (IRS) since 2006.² The whistleblowing literature has tended to focus on the effect on levels of compliance of the presence of potential whistleblowers (Mealem *et al.*, 2010), and the optimal level of incentives for whistleblowing (Yaniv, 2001). Unlike whistleblowing programs, whose aim is to facilitate and incentivize potential informants, the focus of IOVDS is on collecting tax on undisclosed income so as to maximize receipts net of costs.

¹ To overcome this difficulty, the amnesty literature posits that taxpayer's learn new information regarding their own characteristics after the time of the initial reporting decision. In Andreoni (1991), taxpayers learn about their future consumption, and in Malik and Schwab (1991) taxpayers learn about their disutility of tax evasion. Such assumptions are not necessary in the current context.

² According to the IRS Whistleblower Office, an award worth between 15 and 30 percent of the total proceeds that IRS collects can be paid under the scheme if the amount identified by the whistleblower (including taxes, penalties and interest) is more than \$2 million (IRS, 2013).

To date, tax authorities have found various means of acquiring information on offshore holdings. First, in some instances, tax authorities have aggressively exploited legal powers that impel financial organizations to reveal tax-related information. One of the first such IOVDS, the 2007 Offshore Disclosure Facility (ODF), was implemented in the UK following legal action by the tax authority to force five major UK banks to disclose details of the offshore accounts held by their customers. The ODF offered affected taxpayers time-limited access to a reduced 10 percent fine rate if they made a full disclosure. Ireland (2004) and Australia (2009) have also implemented IOVDS following similar legal action.

Second, tax authorities have co-operated with whistleblowers. Following the receipt of information that Swiss bank UBS was actively assisting and facilitating the concealment of taxable income, the IRS launched the 2009 Offshore Voluntary Disclosure Program (OVDP), which offered a time-limited 20 percent penalty to those who made a full voluntary disclosure. This was followed by the 2011 Offshore Voluntary Disclosure Initiative (OVDI), which operated along similar lines, but which targeted clients of other institutions.³ In 2009 the UK implemented two further IOVDS—the New Disclosure Opportunity (NDO) and the Liechtenstein Disclosure Facility (LDF)—in response to, for instance, information acquired on the offshore accounts of around 100 UK citizens from a former employee of a Liechtenstein bank and the receipt from a whistle-blower of information relating to all British clients of HSBC in Jersey (Watt *et al.*, 2012). A list of offshore account holders of HSBC's Geneva branch—seized by French police in 2009—is still the subject of investigation by tax authorities worldwide.⁴ A second list concerning ten offshore jurisdictions is also now in the public domain (see Center for Public Integrity, 2013). To address these and other sources of information, Italy, France, Canada and Hungary are also known to have implemented IOVDS to recover tax on offshore funds (see, e.g., OECD, 2010; Thomson Reuters, 2009).⁵

Third, tax authorities have sometimes exploited changes to legislation. In 2003, for instance, a new European Savings Directive (European Union, 2003) was a part of the impetus for the UK's NDO. Last, tax authorities have taken steps to improve international co-operation through the signing of tax information exchange agreements. For instance, part of the motivation for the UK's LDF was to deal with cases flowing from the signing of a tax information exchange agreement with Liechtenstein in 2009.⁶

IOVDS have proved effective net revenue raisers: the 2009 OVDP in the United States, for instance, raised some USD 3.4 billion (GAO, 2013). In the UK, the 2009 ODF raised nearly £500 million (Treasury Committee, 2012: 14) and cost £6 million to administer (Committee of Public Accounts, 2008: 9)—implying a return of £67 for every £1 spent. This compares favorably with reported yield/cost ratios in the UK of around eight-to-one for traditional audit-based enforcement programs (HMRC, 2006).⁷ Moreover, IOVDS typically raise revenue much faster than does a system relying on (often lengthy) audits.

To our knowledge, however, little systematic is yet understood concerning the optimal design of IOVDS. In particular, we examine how tax authorities should optimally set the rule that determines whether disclosures under an IOVDS are accepted or audited, and the optimal fine rate to apply to accepted disclosures. When the tax authority has information on only some fraction of the clients of a particular financial institution, it faces an additional choice: the set of taxpayers to whom it sends a letter. For instance, prior to the OVDP in the United States, the Swiss authorities agreed to hand the IRS the names of approximately 4,450 United States clients with accounts at UBS. The IRS then had the choice of (i) requiring UBS to write to the 4,450 affected taxpayers informing them that the details of their offshore holding had been handed to the IRS; or (ii) requiring UBS to write to a wider set of its clients (up to the set of all UBS clients with offshore holdings) informing them that the details of their offshore holding *might* have been handed to the

³ See Table 1 and Appendix II of Government Accountability Office (2013) for a full account of the background to, and operation of, these two IOVDS.

⁴ A subset of this list is the so-called Lagarde List—which contains 1,991 names of Greeks with accounts in Switzerland. It was passed to the Greek authorities in 2010 by the then French Finance Minister, Christine Lagarde (Boesler, 2012).

⁵ Rather than implement a bespoke IOVDS, some tax authorities have chosen to handle offshore information through standing mechanisms for voluntary disclosure. In particular, Germany has not to date implemented an IOVDS, but is thought to have raised around €4 billion in voluntary disclosures following the acquisition of data from a Liechtenstein whistleblower (OECD, 2010). Rather than implement a letter campaign, such countries have instead relied on media coverage to inform taxpayers of the information the tax authority has obtained. In this paper, however, we analyse only the optimal design of a bespoke IOVDS in relation to a particular acquisition of offshore information, rather than the optimal design of a general-purpose mechanism for voluntary disclosure.

⁶ Following the signing of tax information exchange agreements with the Isle of Man, Jersey and Guernsey, the UK is now also implementing three identical IOVDS (one for each dependency).

⁷ The ratio of 8:1 is based on the estimated yield/cost ratio for self-assessment non-business enquiry work in 2005-06 of 7.8-to-one.

IRS. In actuality, the IRS chose the second option, and—to prevent taxpayers from inferring whether their information had been handed over—negotiated a confidentiality clause with the Swiss that concealed the criteria by which the 4,450 accounts were selected until after the 2009 OVDP deadline had passed (GAO, 2013). We therefore also examine who tax authorities should communicate with.

In this paper we develop, and analyze with simulation, a two-stage principal agent model of an unanticipated IOVDS. In common with the amnesty models of Andreoni (1991) and Malik and Schwab (1991), we examine both the initial decision by the taxpayer to evade, as well as the taxpayer's subsequent disclosure under the IOVDS. In the first period, taxpayers—who are heterogeneous in respect of initial wealth—decide on a level of offshore evasion (not anticipating a future IOVDS). At the start of the second period, the tax authority acquires information, and commits to the design of an IOVDS.⁸ Taxpayers then make a disclosure (which may be zero) under the IOVDS. Final payoffs are determined when the tax authority responds to each disclosure (by accepting it or by marking it for audit) according to its precommitted rule. Within this setting it is possible to examine a rich range of design questions, including when an IOVDS is a cost-effective enforcement strategy; how generous amnesty incentives should be; which taxpayers should be sent a letter; whether the tax authority should promote strategic ambiguity over the signal it has received; and how it will respond to disclosures.

The principal-agent approach adopted here involves the assumption that the tax authority (principal) can commit to both offering a discounted fine rate to accepted disclosures, and to an auditing rule. If the tax authority cannot commit to the discounted fine rate it communicates to taxpayers it will renege on this rate *ex post*. The IOVDS implemented to date, however, show that tax authorities can commit to discounted fine rates. As in other contexts, one reason this might be the case is that an IOVDS is, in a wider context, only one stage in a repeated game between taxpayers and the tax authority. While renegeing on lower fine rates might be optimal in a one-shot game, the loss of reputation and integrity this action would entail might harm the long-run effectiveness of such IOVDS, and potentially other compliance activities. Regarding commitment to auditing, Stella (1991) notes that most amnesties—even those that are not associated with the receipt of new information—are presaged by claims of increased future enforcement activity. When, however, the tax authority has information on significant amounts of tax liability, this claim seems inherently credible. The tax authority might, nonetheless, be tempted to exaggerate *ex ante* the amount of subsequent auditing it will perform. Commitment here could come from the observation that the record of the tax authority in auditing and prosecuting the taxpayers identified in large information acquisitions is often the subject of close scrutiny by public bodies (see, e.g., GAO, 2013, for the U.S. and Committee of Public Accounts, 2012, for the UK) and the media (e.g., Peev, 2012).

The paper proceeds as follows: Section 2 presents a model of the operation of an IOVDS; Section 3 describes a simulation of the model, and Section 4 presents the main results. Section 5 concludes.

2. Model

In this section we present a two-stage model of the strategic interaction between taxpayers and the tax authority under an IOVDS.

2.1. Offshore Evasion Decision

Let there be a set of taxpayers \mathbf{T} of mass $|\mathbf{T}| = 1$. Each taxpayer in \mathbf{T} possesses an offshore account with a particular bank on which the tax authority will subsequently acquire information. Taxpayers enter period 0, endowed with an exogenous initial wealth w on which the government seeks to levy tax at a marginal rate θ . Initial wealth is distributed across taxpayers according to $g_w(\cdot)$, with mean μ_w , and there exists a continuum of taxpayers at each level of wealth.⁹ We

⁸ The principal-agent approach adopted here involves the assumption that the principal can commit to an audit policy which the agent takes as given. Though important, as in many other contexts, we do not elaborate upon how the principal can create commitments. For detailed discussion of this point, see, e.g., Reinganum and Wilde (1986) and Melamud and Mookherjee (1989). We note that the type of Scheme we examine is based upon those actually implemented in the UK. Empirically, therefore, tax authorities do appear to be able to commit in this way.

⁹ In practice, of course, the set T will be finite, and so all subsets of T are finite too. We explore the case with a continuum of taxpayers as, otherwise, our simulation results are sensitive to the realisations of the random variables contained in the model, making inference uncertain. These difficulties can be overcome in a finite setting by allowing for sufficiently many taxpayers but, in practice, this is feasible only if the taxpayers face a sufficiently simple underlying maximisation problem (ideally with an analytic solution). We, however, explore a relatively complex optimization that must be solved numerically, which limits the feasible number of individual taxpayers we can consider. The infinite setting considered here can be viewed as yielding the long-run average performance of an IOVDS if it were repeated many times.

model an offshore investment as a form of evasion technology that hides income, with probability one, from the tax authority's regular audit program.¹⁰

Taxpayers are assumed to behave as if they (i) possess a utility function $U(\cdot)$; (ii) have monotone ($U' > 0$) and risk averse ($U'' \leq 0$) preferences over wealth; and (iii) maximise expected utility. As, for many years, offshore accounts were considered to be entirely untraceable by national tax authorities, we assume that the IOVDS that will subsequently occur in period 1 is wholly unanticipated by the taxpayer at time zero. Accordingly, a w -taxpayer's problem in period 0 may be written as

$$\max_{E_w \in [0, w]} U([1 - \theta][w - E_w]) + \delta U([1 - \theta][w - E_w] + E[1 + r]).$$

An interior optimum is characterized by the first order condition

$$\frac{\delta[r + \theta]}{1 - \theta} = \frac{U'([1 - \theta][w - E_w])}{U'([1 - \theta][w - E_w] + E[1 + r])}. \quad (1)$$

The distribution of offshore evasion implied by condition (1) we write as $g_{\mathbf{T}}(E_w)$.

2.2. IOVDS

2.2.1. Information and Communication

At the beginning of period 1 the tax authority requires a given financial institution to hand over information concerning the offshore holdings of a subset, $\mathbf{I} \subseteq \mathbf{T}$, of the affected bank's clients. We write the distribution of \mathbf{I} as $g_{\mathbf{I}}(E_w) = \phi g_{\mathbf{T}}(E_w)$, where $\phi \in [0, 1]$ is the proportion of the taxpayers in \mathbf{T} that belong to at each level of wealth. We assume that the tax authority knows ϕ , so, having observed $g_{\mathbf{I}}(\cdot)$, it can infer $g_{\mathbf{T}}(\cdot)$ as $\phi^{-1} g_{\mathbf{I}}(\cdot)$. The tax authority may incur costs in acquiring information. In several of the IOVDS discussed in the Introduction, however, the tax authority acquired information at essentially zero-cost. In other cases payments to whistleblowers were made, but the amounts involved (to be the extent these are observed) appear relatively modest in relation to the revenue generated.¹¹ As the optimal reward to whistleblowing is the subject of its own dedicated literature (e.g., Yaniv, 2001), we assume that the tax authority acquires its information at zero cost, so as to focus on other aspects of the analysis.

The tax authority's signal need not permit a fully accurate assessment of the undisclosed tax liability. Whistleblower data, for instance, could be inaccurate or incomplete; especially as the whistle-blowers motives for blowing the whistle are not known. We wish to emphasize, however, that tax liabilities remain uncertain even when the tax authority observes perfectly the amount of a taxpayer's offshore assets (such as when this information is provided directly by a financial institution). At one extreme, identified taxpayers may be fully compliant, for possessing an offshore account is not, in itself, illegal, if properly declared to the tax authority. Alternatively, the offshore assets may have been properly taxed at source, making the taxpayer liable only for undeclared interest and capital gains on the assets. The largest liabilities occur if the offshore assets themselves are liable for, e.g., inheritance tax.

Accordingly, we suppose that the tax authority's private signal is noisy (but unbiased).¹² In particular, the signals are generated according to $s_i = q_i \theta E_w$, where q_i is drawn for each taxpayer from a random variate \tilde{q} with cumulative distribution function (cdf) $G_q(\cdot)$, mean one, and support $(\underline{q}, \bar{q}) \in \mathbf{R}_+$. We treat the variance of \tilde{q} —denoted σ_q^2 —as measuring the quality of the tax authority's signal. We shall assume that both the tax authority and taxpayers observe \tilde{q} but neither observes the realization q_i for any individual taxpayer.

¹⁰ A cost to offshore evasion, as allowed for in, e.g., Lee (2001), can readily be added to the model with predictable consequences—higher evasion costs lower optimal evasion. We assume a zero cost therefore.

¹¹ The UBS employee who acted as an IRS informer allegedly received a payment of USD 104 million, but in the context of some USD 3.4 billion that was eventually raised by the resulting IOVDS (GAO, 2013). The British tax authority is reported to have paid a former Liechtenstein bank employee a fee of just £100,000 for information regarding more than £100 million of offshore funds (Oates, 2008). The German tax authority is also understood to have paid an undisclosed sum to the same individual for information regarding the accounts of its citizens.

¹² If the tax authority knows its signal to be biased, it may simply inflate its observed signals by the average degree of bias to obtain unbiased estimates. Hence, we do not investigate this possibility.

Having observed the signal, the tax authority chooses a letter set, $\mathbf{L} \supseteq \mathbf{I}$. It then requires the affected financial institution to write to all clients belonging to \mathbf{L} outlining the terms of an IOVDS. The distribution function of \mathbf{L} we denote by $g_{\mathbf{L}}(E_w)$ and its cdf by $G_{\mathbf{L}}(E_w)$. To investigate the choice of \mathbf{L} we specify the density of \mathbf{L} as $g_{\mathbf{L}}(\cdot) = g_{\mathbf{I}}(\cdot) + \kappa[g_{\mathbf{T}}(\cdot) - g_{\mathbf{I}}(\cdot)]$, where $\kappa \in [0, 1]$ is a choice parameter of the tax authority. At one extreme, if $\kappa = 0$, then $g_{\mathbf{L}}(\cdot) = g_{\mathbf{I}}(\cdot)$, which we interpret as every member of \mathbf{L} belonging to \mathbf{I} . At the opposite extreme, if $\kappa = 1$ then $g_{\mathbf{L}}(\cdot) = g_{\mathbf{T}}(\cdot)$ in which case all taxpayers in \mathbf{T} receive a letter ($\mathbf{L} = \mathbf{T}$). For intermediate values $\kappa \in (0, 1)$ a fraction κ of the taxpayers in $\mathbf{T} \setminus \mathbf{I}$ at each level of wealth are sent a letter (in addition to all members of \mathbf{I}). Note that if $\mathbf{I} = \mathbf{T}$ then necessarily $\mathbf{L} = \mathbf{T}$, so any choice of κ is weakly optimal.

2.2.2. Responses, Penalties and Costs

The letter invites taxpayers to make a voluntary disclosure of an amount of offshore income. The amount the taxpayer chooses to disclose we denote by $x \in [0, E]$, and we treat choosing not to make a disclosure as equivalent to making the disclosure $x = 0$. In connection with this disclosure, the tax authority commits to an *audit rule* and a fine rate for accepted disclosures. Based on the design of the NDO in the UK (see HMRC, 2009), we assume the audit rule takes the following form: (i) if $i \in \mathbf{I}$ and $s_i - \theta x_i \leq a$ then accept the disclosure (state A); (ii) if $i \in \mathbf{I}$ and $s_i - \theta x_i > a$ perform an audit (state H); (iii) if $i \notin \mathbf{I}$ —as will be the case for taxpayers in $\mathbf{L} \setminus \mathbf{I}$ —then accept the disclosure.¹³ The parameter a , the value of which is assumed not to be communicated in the letter, determines the leniency of the tax authority in accepting disclosures.

Let $f_H > 0$ be the applicable regular, or unincentivized, fine rate prescribed in national legislation, that we assume the tax authority takes as given. Offshore evasion revealed at audit that was not disclosed under the Scheme is assumed to be subject to this rate. If a disclosure is accepted, however, the taxpayer faces a fine—the level of which is chosen by the tax authority—on the implied tax liability at an incentivized rate, $f_A \in [0, f_H]$. The assumption that only f_A is actively chosen by the tax authority is in one sense a simplification for reasons of tractability, but, in another, reflects practice in the UK, where IOVDS been designed, wherever possible, to fit into existing penalty legislation so as to minimize, or mitigate altogether, the need for new legislation.

The two response states entail different administration costs for the tax authority. If the tax authority accepts a disclosure, we assume there are no further administration costs (in the UK, taxpayers were told to include a fine payment, at rate f_A , in their disclosure). If, however, the tax authority audits the disclosure it incurs a per-disclosure cost of $c_H > 0$.¹⁴

2.3. Optimal Disclosure

Taxpayers enter period 1 with an (endogenous) initial post-tax wealth of $W_w = w(1 - \theta) + \theta E$, of which an amount E is hidden offshore. The threshold parameter a is not known to taxpayers. We suppose, however, that each taxpayer holds a common belief, \tilde{a} , which is a probability distribution $g_a(\cdot)$ over a , with mean μ_a and variance σ_a^2 on the support $[\underline{a}, \bar{a}]$.

Taxpayers who receive a letter must assess the likelihood that the tax authority has observed a signal, $p_{\mathbf{I}\mathbf{L}}$. We assume that taxpayers know the size of both \mathbf{I} and \mathbf{L} (but not the identities of the taxpayers belonging to either set), in which case $p_{\mathbf{I}\mathbf{L}} = |\mathbf{I}|/|\mathbf{L}|$. The subjective expected utility of a w -taxpayer who receives a letter is then given by

¹³ The assumption that disclosures from taxpayers in $\mathbf{L} \setminus \mathbf{I}$ are always accepted (and that taxpayers know this) is clearly a simplification. The tax authority might, in practice, want to follow up positive disclosures with audits.

¹⁴ The assumption that the tax authority must perform an audit for legal verification purposes (even if it knows the true liability with certainty) is standard in the optimal auditing literature (see, e.g., Reinganum and Wilde, 1985; Morton, 1993).

$$\begin{aligned}
EJ(x_w) &= [1 - p_{\mathbf{IL}}]U(W_A(x_w)) \\
&+ p_{\mathbf{IL}} \int G_q \left(\frac{\varphi + \theta x_w}{\theta E_w} \right) U(W_A(x_w)) dG_a(\varphi) \\
&+ p_{\mathbf{IL}} \int \left[1 - G_q \left(\frac{\varphi + \theta x_w}{\theta E_w} \right) \right] U(W_H(x_w)) dG_a(\varphi)
\end{aligned} \tag{2}$$

where $W_j(x)$ is the payoff in state j :

$$\begin{aligned}
W_A(x_w) &= W_w - \theta[1 + f_A]x_w; \\
W_H(x_w) &= W_A - \theta[1 + f_H][E_w - x_w].
\end{aligned}$$

Each w -taxpayer's problem is to choose an income disclosure $x_w \in [0, E_w]$ to maximize expected utility, subject to the equilibrium consistency condition

$$\mu_a = a, \tag{3}$$

which implies that taxpayers' beliefs over the threshold a are centered on the true value. The first order condition for the taxpayer implicitly defines a *disclosure function*, $x_w = d(E_w)$, which maps each level of offshore evasion to an optimal disclosure.

2.4. Optimal Enforcement

The tax authority commits to a choice of the set of parameters $\{a, f_A, \kappa\}$ before taxpayers choose x . Owing to the stochastic nature of q_i , two identical agents who evade the same amount, and make the same disclosure, may nevertheless experience different response states. For a given choice of $\{a, f_A, \kappa\}$ the proportion of w -taxpayers belonging to \mathbf{I} that will experience response state j is given by

$$\begin{aligned}
p_{wA}(E_w) &= G_q \left(\frac{a + \theta d(E_w)}{\theta E_w} \right); \\
p_{wH}(E_w) &= 1 - p_{wA}(E_w);
\end{aligned}$$

and receipts r_j from a w -taxpayer in state j are given by

$$r_{wj}(E_w) = \begin{cases} \theta[1 + f_A]d(E_w) & \text{if } j = A; \text{ and} \\ r_{wA}(E_w) + \theta[1 + f_H][E_w - d(E_w)] & \text{if } j = H. \end{cases}$$

Hence, total revenue—which, by the law of large numbers, is certain—is written as

$$r_w(E_w) = \begin{cases} \theta[1 + f_A]d(E_w) & \text{if } j = A; \text{ and} \\ r_{wA}(E_w) + \theta[1 + f_H][E_w - d(E_w)] & \text{if } j = H; \end{cases} \tag{4}$$

and total costs are

$$C = p_{\mathbf{IL}}c_H \int p_{wH}(E_w) dG_{\mathbf{L}}(E_w).$$

The tax authority's problem is then to choose $\{a, f_A, \kappa\}$ to maximize revenue net of costs, $R - C$, given taxpayers' optimal choice of x_w and $f_A \in [0, f_H]$. As is standard, we also assume that the tax authority cannot fine a taxpayer by an amount that exceeds their wealth, such that, for all w , $W_H(x_w) \geq 0$.¹⁵

3. Simulation

In this section we detail a version of the model of section 2 that we shall go on to simulate in section 4. We begin by specifying functional forms for taxpayer utility (power), evasion cost (quadratic), and the distribution of initial wealth (bounded Pareto):

$$U(c; \gamma) = \frac{c^{1-\gamma} - 1}{1-\gamma}; \quad (5)$$

$$\rho(E, w) = \beta \left[\frac{E}{w-E} \right]^2; \quad (6)$$

$$g_w(w) = \frac{\alpha w^\alpha \bar{w}^{-\alpha-1}}{1 - \left[\frac{\underline{w}}{\bar{w}} \right]^\alpha}; \quad (7)$$

where $\gamma \geq 0$, $\alpha > 0$, $\beta > 0$, and $\{\underline{w}, \bar{w}\}$ are the minimum and maximum of observed initial wealth. In period 0, when taxpayers make a riskless choice, the parameter γ is best interpreted as the elasticity of the marginal utility of wealth. When, however, in period 1, the taxpayer makes a decision involving risk, γ is more conventionally interpreted as the Arrow-Pratt coefficient of relative risk aversion.¹⁶ The parameter β varies the marginal cost of evasion. The Pareto distribution provides a close fit to the far right tail of the wealth distribution (see, e.g., Coelho *et al.*, 2008; Klass *et al.*, 2006), which is disproportionately where the taxpayers affected by IOVDS are located.¹⁷ The parameter α is a shape parameter that we use to vary the degree of skewness.

The noise in the tax authority's signal is assumed to be a truncated normal distribution on the support $[0, 2]$. The public signal \tilde{a} is assumed to be a discrete distribution with non-zero mass at the values $\{a_k\}_{k=1}^3 = \{\mu_a [1 + \delta \kappa]\}_{k=1}^3$, where $\delta \in (0, 1)$. The probabilities assigned to each a_k are taken from a normal distribution centered on μ_a :

$$\left\{ \frac{N[a_k; \mu_a, \sigma_a]}{\sum_{m=1}^3 N[a_m; \mu_a, \sigma_a]} \right\}_{k=1}^3.$$

The results we present are for the following set of parameter values:

$$\begin{aligned} \alpha &\approx 0; \beta = 0.04; c_H = 1.00; \delta = 0.80; f_H = 1.75; \gamma = 2.00; \\ \phi &= 1.00; \sigma_a = 1.00; \sigma_q = 0.45; \theta = 0.30; \underline{w} = 5.00; \bar{w} = 15.0. \end{aligned} \quad (8)$$

We term the parameter values in (8) as the baseline values. When we wish to examine the effects of varying the value of an individual parameter while holding the remaining parameters constant, we use these baseline values to

¹⁵ We focus on specifications of this problem that give a unique optimum for a . This rules out two degenerate versions of the model, one (low a) in which all taxpayers know ex-ante that their disclosure will be audited, whatever the realisation of \tilde{q} , and another (high a) in which all taxpayers know ex-ante that their disclosure will be accepted, whatever the realisation of \tilde{q} . In both these cases the taxpayer's expected utility in (2) becomes independent of μ_a (and so, in equilibrium, also a).

¹⁶ Other specifications of utility such as constant absolute risk aversion or mean-variance utility can instead be used and yield similar results. However, the assumption of constant absolute relative risk aversion has stronger empirical support (Wakker, 2008).

¹⁷ For the purposes of simulation, we employ a discrete approximation to $g_w(\cdot)$. In particular, the discrete distribution places a mass n^{-1} on each of n values, $\{[x_w^{-1}(\underline{w} + [\bar{w} - \underline{w}]/n - 1)^\alpha]^\alpha\}_{k=1}^n$.

specify the remaining parameters. The unincentivized fine rate f_H reflects the 2009 OVDP, in which U.S. taxpayers in state H were liable, in addition to fines payable in state A , for a fraud penalty of 75 percent of the unpaid tax. In the UK, civil fraud legislation sets $f_H \in [0.3, 1]$, the lower bound applying if noncompliance is judged to be through careless error (see, e.g., HMRC, 2012). Most empirical estimates of the coefficient of relative risk aversion are in the neighborhood of two (Meyer and Meyer, 2005), so we adopt this value for γ . We set ϕ to unity (which implies $I = T$) as this choice enables us to examine the optimal choice of a and f_A separately from the choice of the letter set. That is, when ϕ takes its baseline value, any κ is weakly optimal, so this parameter effectively drops out of the model. When, however, we lower ϕ below the baseline value we are then able to examine the optimal size of the letter set. The shape parameter α in the Pareto distribution is set arbitrarily close to zero: this choice minimizes the skewness of initial wealth, such that, when we raise α above zero, we can then investigate the effects on the optimum as skewness increases.

We choose the remaining parameter values such that the model predicts an interior optimum for both a and f_A ; and where the optimum is broadly in line with the oral testimony of HMRC officials in Committee of Public Accounts (2012), which stated that the overwhelming majority of disclosures under the ODF and NDO Schemes were accepted. The specification in (8) is not necessarily unique in meeting these desiderata, yet, having simulated many specifications of the model, we find no evidence that our qualitative results are sensitive to the precise choice of the baseline parameters: qualitatively similar results to those we present here are obtained for a wide range of plausible specifications consistent with an interior maximum for both a and f_A .¹⁸

We locate the solution to the taxpayer's problems in periods 0 and 1 using a simple direct search (compass) algorithm.¹⁹ For the government's problem, which involves a potentially more complex objective function, we first perform a fine grid search over (a, f_A, κ) space to find the region of the solution, before using compass search to locate the solution precisely.²⁰ At the equilibrium associated with the baseline parameter values in (8) we obtain $a = 1.04$ and $f_A = 0.59$, which implies that taxpayers whose disclosure are accepted receive an 22 percent reduction relative to the unincentivized fine rate. At these equilibrium values the proportion of taxpayers that have their disclosure accepted is decreasing in wealth: no w -taxpayers have their disclosure audited, up to around seven percent for \bar{w} -taxpayers. Offshore holdings are increasing absolutely in initial wealth, but declining as a fraction of initial wealth: the latter measure varies between 32 percent for w -taxpayers, and 21 percent for \bar{w} -taxpayers. Disclosures are found to be increasing absolutely and relative to offshore assets, where again the latter measure varies between 41 percent for w -taxpayers, and 59 percent for \bar{w} -taxpayers.

4. Analysis

In this section we analyze the version of the model set out in section 3 using analytical and simulation approaches.

4.1. Offshore Holdings

Although the full model is too complex to admit an analytic treatment, the taxpayer's choice of offshore investment amount in period 0 is sufficiently simple to be theoretically tractable. Although this choice is not our principal interest, an understanding of the separate properties of each choice made by taxpayers is instructive for interpreting the later results for optimal enforcement. We have the following Proposition on the comparative statics properties of the optimal level of offshore assets (the proof being in the Appendix):

Proposition 1 For $\gamma \in [0, 2]$ and $\underline{w} \geq [1 - \theta]^{-1}$ the following hold for a taxpayer at an interior maximum:

$$\frac{\partial E}{\partial \beta} < 0; \quad \frac{\partial E}{\partial \gamma} < 0; \quad \frac{\partial E}{\partial \theta} > 0; \quad \frac{\partial E}{\partial w} > 0.$$

¹⁸ Quantitatively, however, the locations of various kink- and break-points in the model are sensitive to the precise parameter values used. Given the uncertainty over the true empirical values of some of our parameters, we would therefore caution against interpreting our results in a strict quantitative sense.

¹⁹ For a description of the compass search algorithm and a wider review of direct search methods see, e.g., Kolda *et al.* (2003). We employ the method in Lewis *et al.* (2007) when searching close to one or more parameter boundaries.

²⁰ At this solution we obtain an average for the ratio for the total disclosures to total hidden income of 0.54, and with just over 95 percent of disclosures accepted.

According to Proposition 1, wealthier individuals place more assets offshore, which is consistent with the prevalence of high net worth individuals within the lists released by whistleblowers. Higher marginal tax rates provide a greater incentive to move wealth offshore, but higher costs act as a disincentive. The comparative static effect for γ is more complex, but under the two conditions in the Proposition—both of which are satisfied by the baseline specification in (8)—can be shown to be negative. Note that, as we model the period 0 problem as a choice under certainty, γ cannot be formally interpreted as a risk parameter in this context. Proposition 1 shows that higher values of γ lower the optimal offshore investment, just as an increase in risk aversion reduces evasion in the canonical evasion decision under uncertainty (e.g., Allingham and Sandmo, 1972). The offshore investment is independent of the remaining model parameters that do not appear in Proposition 1.

4.2. Optimal Disclosure

We now characterize the period 1 problem of taxpayer's choice of a disclosure $x \geq 0$ for a given choice of the set $\{a, f_A, \kappa\}$ by the tax authority. When $\{a, f_A, \kappa\}$ are chosen endogenously, the disclosure is affected indirectly by every parameter of the model, but it is instructive to distinguish between (i) the set of parameters that affect the disclosure through this indirect route only— $\{a, \alpha, c_H, \mu_w\}$ —and (ii) the sets of parameters that affect the disclosure: (1) directly— $\{f_A, f_H, \kappa, \mu_a, \sigma_a, \sigma_q\}$; (2) indirectly through the endogenous determination of $E - \{\beta\}$; and (3) through both the routes in (1) and (2)— $\{\gamma, \theta, w\}$. Taking these sets in turn, note that the two parameters we use to characterize the distribution of initial wealth belong to the set in part (i), for only a taxpayer's own initial wealth enters the expected utility function in (2). The parameters a and c_H also belong to this set, the former as it enters the optimal disclosure problem only through the equilibrium consistency condition in (3), and the latter as it affects only net revenue.

The majority of parameters directly affect the optimal disclosure. The comparative statics for some elements of this set can be shown analytically to be unambiguously of a given sign.²¹ First, as the fine rates $\{f_A, f_H\}$ enter equation (2) only in the payoffs, it is possible to show that these parameters have opposite effects: an increase in the incentivized fine rate decreases the optimal declaration ($\partial x / \partial f_A < 0$), but an increase in the unincentivized fine rate increases the optimal disclosure ($\partial x / \partial f_H > 0$). It is also possible to show that an increase in the taxpayer's mean belief, μ_a , must reduce the optimal disclosure ($\partial x / \partial \mu_a < 0$), for it unambiguously reduces the subjective probability a taxpayer assigns to the event of their disclosure being audited. Taking together the findings for f_A and μ_a , it follows that *a tax authority seeking to maximize disclosures under an IOVDS always chooses a and f_A to attain their respective lower bounds*. A net-revenue maximizing tax authority might, however, might set both a and f_A parameters above their lower bounds, in the former case as it must take into account the cost of performing audits, and in the latter as fine payments contribute positively to revenue.²²

The parameter κ enters expected utility only through $p_{\mathbf{L}}$, which itself enters expected utility linearly. Using the definitions in section 2, we obtain

$$p_{\mathbf{L}} \equiv \frac{|\mathbf{I}|}{|\mathbf{L}|} = \frac{|\mathbf{I}|}{|\mathbf{I}| + \kappa[1 - |\mathbf{I}|]},$$

which implies $\partial p_{\mathbf{L}} / \partial \kappa < 0$. As it is also possible to show that $\partial x / \partial p_{\mathbf{L}} > 0$, these results together imply that an increase in κ lowers the optimal disclosure.²³ The fundamental trade-off facing the tax authority when choosing the size of the letter set is therefore between the marginal revenue from receiving a disclosure from one extra taxpayer (who would not otherwise have made a disclosure under the Scheme); and the marginal revenue loss from a reduction in the disclosures of all existing members of \mathbf{L} (from a fall in $p_{\mathbf{L}}$). The signs of the comparative statics effects for the two remaining variables that directly affect the optimal disclosure (σ_a and σ_q) are much more difficult to ascertain

²¹ The proofs for the analytical results given in this section use the same techniques as used in the proof of Proposition 1. We do not give them here, however, as the workings are burdensome, and offer little or no insight beyond what is readily intuitive.

²² Similarly, Rablen (2014) shows that a net-revenue maximizing tax authority does not maximize voluntary compliance.

²³ To see the second result intuitively, observe that, if $p_{\mathbf{L}} = 0$ (such that taxpayers give no weight to belonging to \mathbf{I}) then a disclosure $x = 0$ is always optimal. As $p_{\mathbf{L}}$ is increased, the subjective probability that a taxpayer assigns to the event that their disclosure is accepted unambiguously falls, which pushes up the optimal disclosure.

analytically. We therefore numerically compute the sign of these effects when the exogenous parameters of the model are set according to the baseline values in (8) and the endogenous parameters take their equilibrium values. We use the subscript N to distinguish numerical derivatives from those determined analytically.²⁴ We find that taxpayer's disclose more as their uncertainty over the true value of a increases, but less as the quality of the tax authority's signal worsens: $\partial x / \partial \sigma_a >_N 0$; $\partial x / \partial \sigma_q <_N 0$.

The cost of offshore evasion is the only parameter that directly affects the offshore investment decision, but not the disclosure decision. From Proposition 1 we have $\partial E / \partial \beta < 0$, and we numerically obtain $\partial x / \partial E >_N 0$, which (by the chain rule) implies $\partial x / \partial \beta < 0$. As β lowers both offshore evasion and the amount disclosed, what is the overall effect on the undisclosed liability, $E - x$? We find (numerically) that evasion is more sensitive to β than is the disclosure, such that the undeclared liability is decreasing in evasion costs. The comparative statics effects for the remaining three parameters—those that affect the optimal disclosure both directly and through E —cannot be unambiguously determined analytically. Using the numerical procedure described above, we obtain:

$$\frac{\partial x}{\partial \gamma} <_N 0; \quad \frac{\partial x}{\partial \theta} >_N 0; \quad \frac{\partial x}{\partial w} >_N 0.$$

At first blush, the finding that an increase in risk aversion lowers the optimal disclosure appears counter-intuitive. It arises as an increase in γ lowers the offshore evasion in period 0, such that, although x falls absolutely, it represents a larger proportion of the offshore investment.²⁵ An increase in θ increases the optimal disclosure, both directly and through its effect on E . Although there are competing effects (θ increases both evasion and the disclosure) we find the undisclosed liability, $E - x$, to be increasing in the tax rate. Last, we find that the optimal disclosure is increasing in wealth, but the direct and indirect effects that make up this result go in opposite directions. Holding E fixed, an increase in initial wealth induces a taxpayer to disclose *less*—a result that follows from the property of decreasing absolute relative risk aversion of the utility function in (5). On the hand, wealthier taxpayers, choose a higher E in period 0, which induces them to disclose *more*. Once again, as w increases both offshore evasion and the amount disclosed, its effect on the undisclosed liability is unclear. We find, however, that the amount evaded increases faster with w than does the amount disclosed, such that $E - x$ is an increasing function of initial wealth.

4.3. Optimal Enforcement

We now examine the tax authority's optimal choice of the parameter set $\{a, f_A, \kappa\}$. As the analysis of these parameters is much too complex for analytic methods to be tractable, we now rely wholly on the findings of the simulation. Figure 1 presents the results of an exercise in which we vary each individual parameter across an interval of values, holding all other exogenous variables at their baseline value. The panels are presented in pairs: the left panel of each pair shows the results for the optimal a and f_A , and the implied optimal proportion of taxpayers who have their disclosure accepted, $|A| \in [0, 1]$. The right panel of each pair shows the tax authority's net revenue ($R - C$) and the measures

$$\mu_E \equiv \frac{\int E_w \, dG_L(E_w)}{\int dG_w(w)}, \quad \mu_x \equiv \frac{\int d(E_w) \, dG_L(E_w)}{\int E_w \, dG_L(E_w)},$$

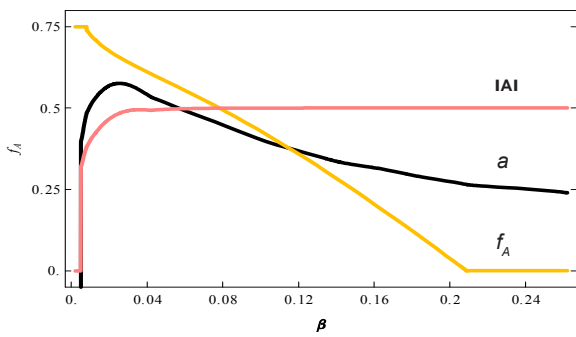
where $\mu_E \in [0, 1]$ is an aggregate measure of the fraction of initial wealth invested offshore by the members of the letter set, and μ_x is an analogous measure for the fraction of offshore evasion disclosed.

Before considering the individual results in detail, we remark on two features of the optimal determination of the incentivised fine rate that feature in several of the results: in panels {g, i, k, m, q} f_A is observed to hit its upper bound, $f_A = f_H$, while in panels {a, o, q} it is observed to hit its lower bound $f_A = 0$.

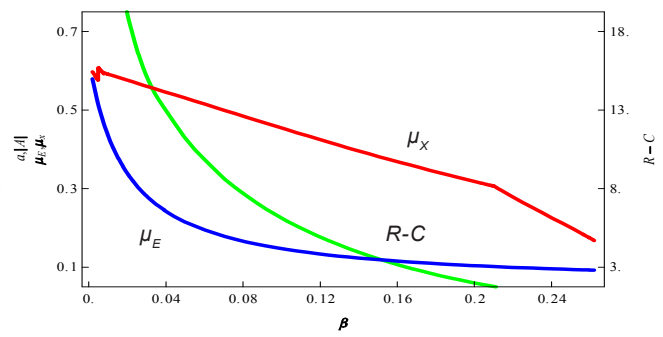
²⁴ We stress that, whereas the analytical derivatives hold in general, the numerical derivatives are for the particular specification of the model in (8). Although we find the signs of the numerical effects we examine to be robust across many specifications, they cannot be assumed to hold in general.

²⁵ Accordingly, if E is held constant, then indeed we find that an increase in γ increases the optimal disclosure.

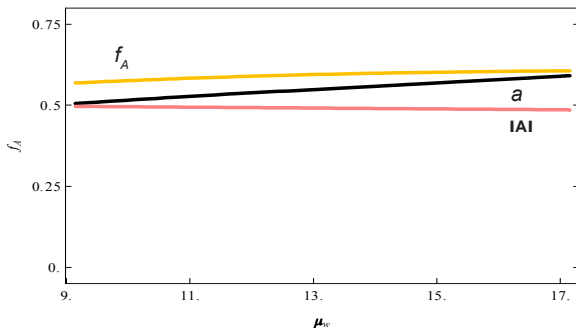
FIGURE 1. Effect of parameter shift on f_A (yellow); a (black); $|A|$ (pink); μ_E (blue); μ_X (red); and $R-C$ (green).



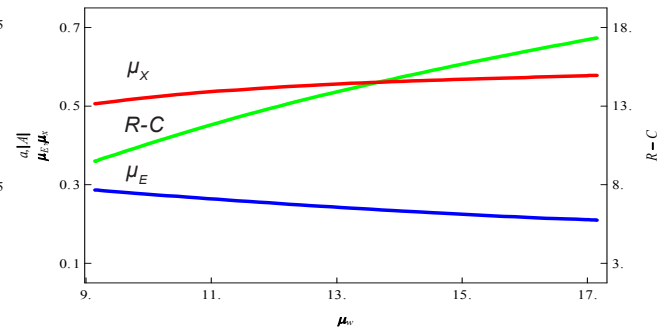
(a) Cost of offshore evasion



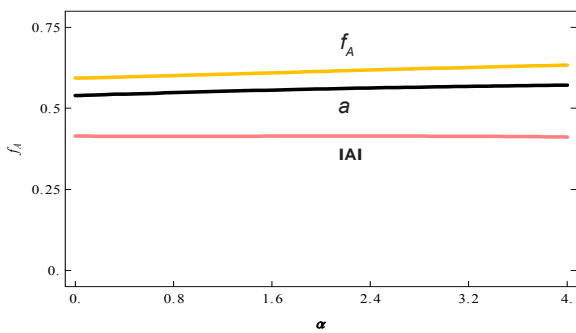
(b) Cost of offshore evasion



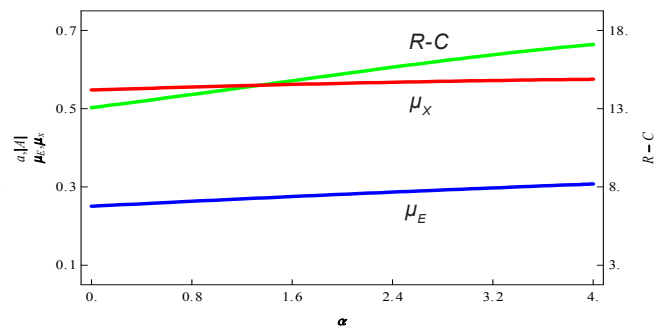
(c) Mean initial wealth



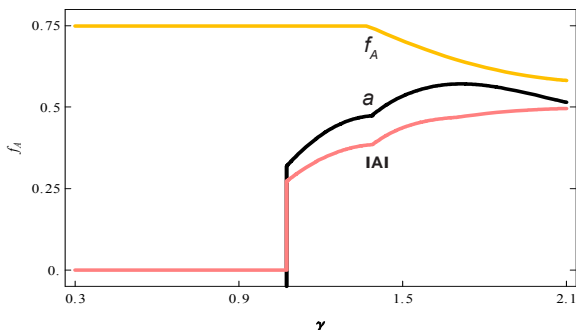
(d) Mean initial wealth



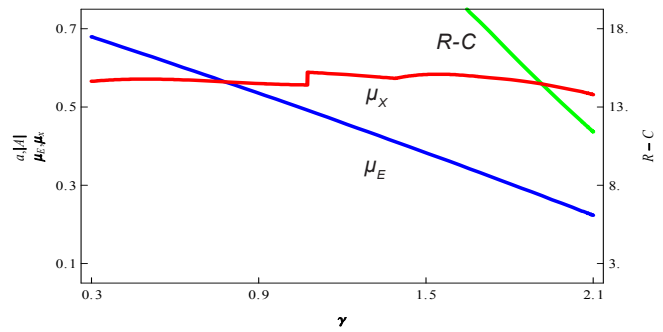
(e) Skewness of initial wealth



(f) Skewness of initial wealth



(g) Curvature of utility



(h) Curvature of utility

FIGURE 1 (Continued). Effect of parameter shift on f_A (yellow); a (black); $|A|$ (pink); μ_E (blue); μ_X (red); and $R-C$ (green).

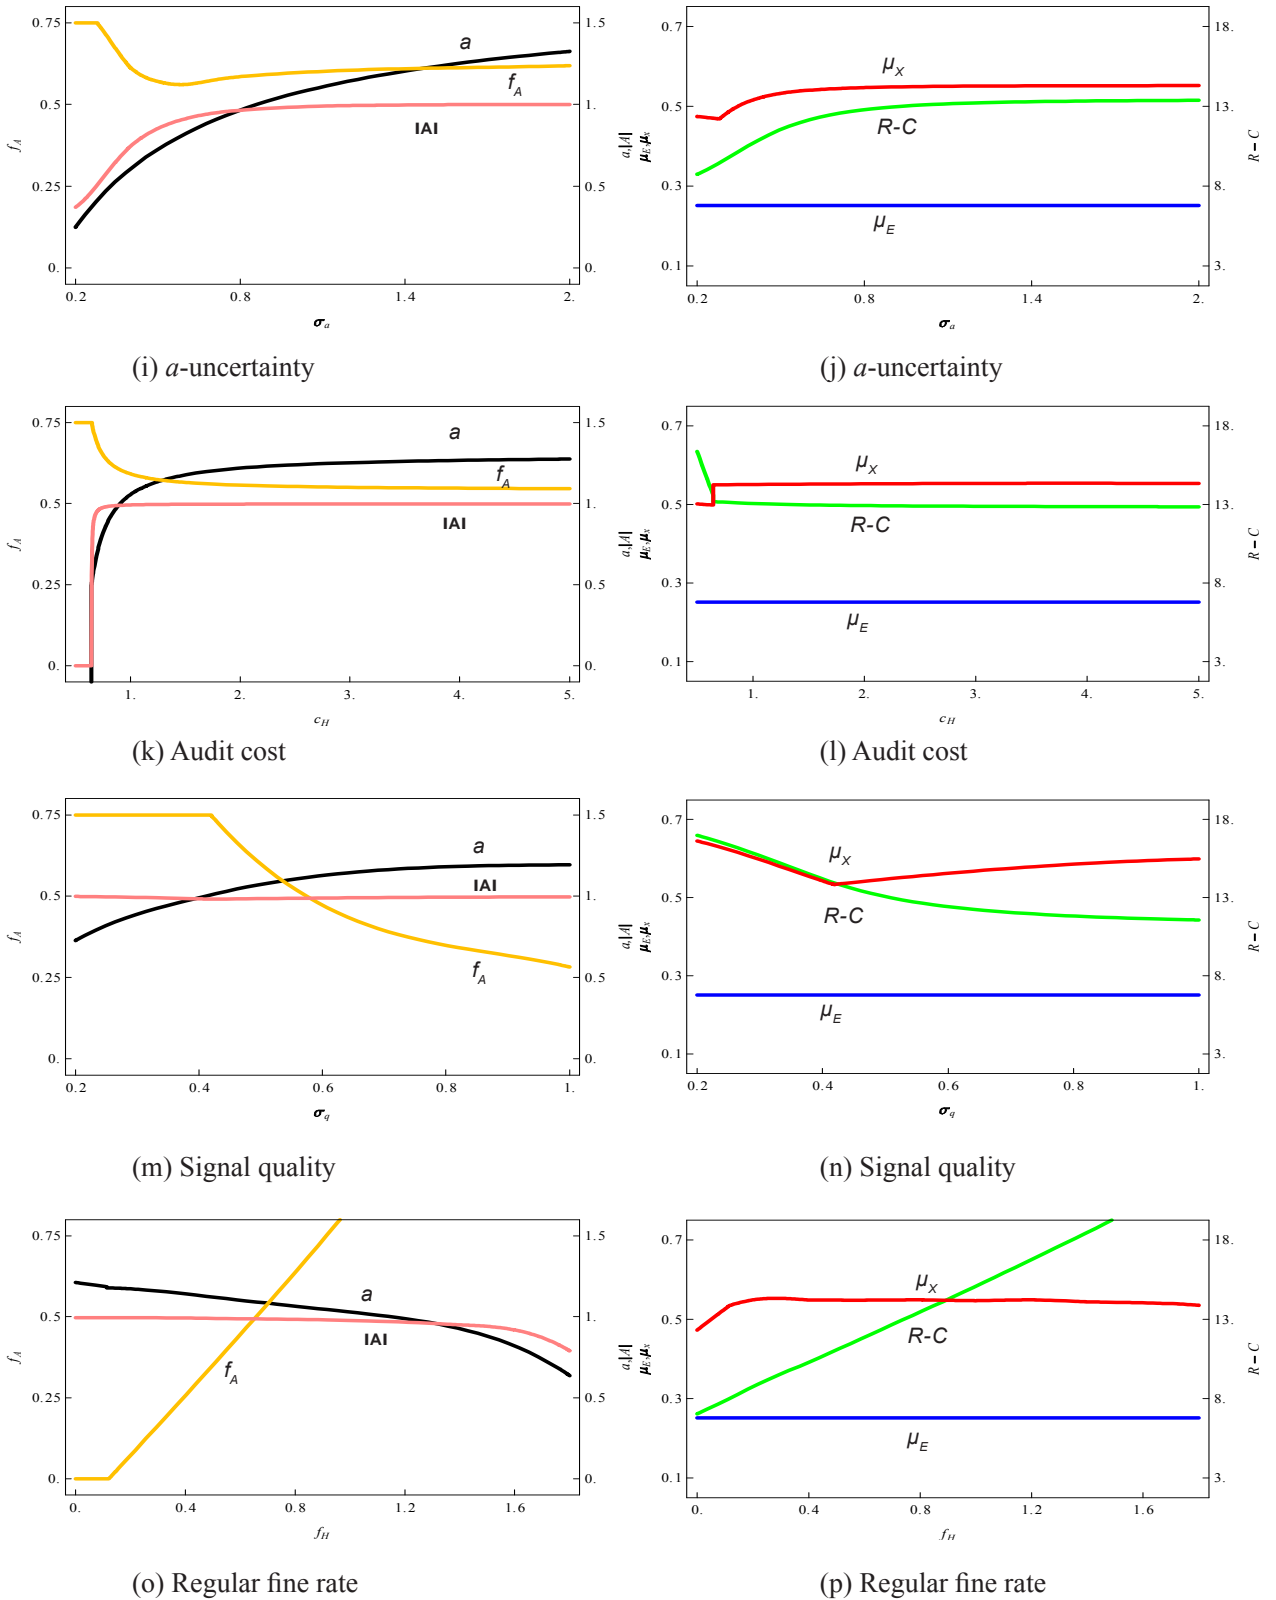
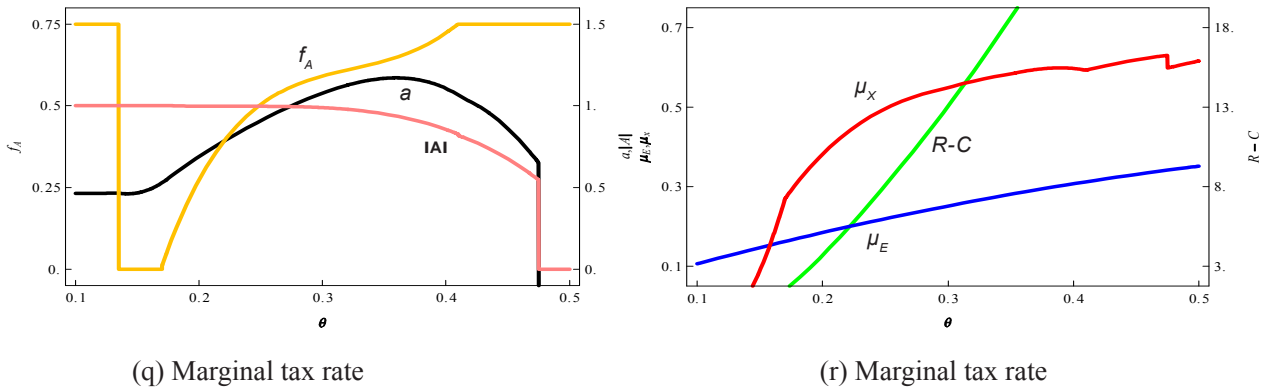


FIGURE 1 (Continued). Effect of parameter shift on f_A (yellow); a (black); $|A|$ (pink); μ_E (blue); μ_X (red); and $R-C$ (green).



REMARK 1 A net-revenue maximizing tax authority may optimally wish to set

- (i) $f_A > f_H$;
- (ii) $f_A < 0$.

Part (i) of Remark 1 may at first seem perverse, for it apparently rewards dishonesty. To understand this finding, note that if a disclosure is accepted the taxpayer pays a fine on the amount $x \leq E$, whereas if a disclosure is audited the taxpayer pays a fine on E in full. Hence, for $x < E$, a taxpayer might still face a smaller fine if their disclosure is accepted even when $f_A > f_H$. Recognizing this, the tax authority can levy a higher fine on accepted disclosures without all taxpayers migrating to state H . Part (ii) of the remark is that a tax authority may not simply wish to offer reduced penalties—the situation considered in the tax amnesty literature—but to forego some of the tax owed by applying a rate of tax less than θ to disclosures under the Scheme. This eventuality arises in the model when the cost of audits outweighs their expected return, such that the tax authority rationally wishes to perform very few audits, or none at all. In these circumstances the use of extreme incentives for honesty can be optimal.

4.3.1. Taxpayer variables

We now turn to the individual results in detail. Panels (a)-(j) of Figure 1 show results for the set of variables that describe characteristics of the taxpayers: $\{\alpha, \beta, \gamma, \mu_w, \sigma_a\}$. We begin with the parameter β that regulates the cost of offshore evasion, the results for which are shown in panels (a) and (b). For very low β , undisclosed liabilities are significant, implying high returns to auditing. In this case the tax authority audits all disclosures, which is performed in a manner consistent with net-revenue maximization by setting a to its lower bound and by offering no incentive for honesty. Conversely, at very high β , undisclosed liabilities become so small that audit activity is undesirable. In this case the tax authority audits no disclosures, which is performed in a manner consistent with net-revenue maximization by offering maximal incentives for honesty, and by being maximally lenient. Over the an intermediate range of β (the range most empirically plausible) the tax authority responds to the decline in undisclosed liabilities as β increases by offering *stronger incentives* for honesty, and by being *more lenient* over accepting disclosures. Note that the tax authority becomes less lenient in an absolute sense (a declines) once the majority of disclosures are being accepted. Note however, that even when the tax authority does become less lenient, it nonetheless does so to a sufficiently small extent that the proportion of disclosures that are accepted continues to increase (panel (a)).

How does the distribution of initial wealth affect optimal enforcement? Beginning with the effects of the mean, panels (c) and (d) of Figure 1 show the results of an exercise in which we simultaneously increase \underline{w} and \bar{w} . We see that increases in μ_w are associated with decreasing incentivization of honesty but increasing leniency. The tax authority performs gradually fewer audits, yet disclosures drift upwards, as does net revenue. The size of these effects is observed to be quite modest, which is a consequence of the observation in section 4.2 that the distribution of initial wealth has only an indirect impact on the optimal disclosure via its effects upon $\{a, f_A, \kappa\}$.

Published data from the 2009 OVDP in the U.S. shows that the mean assessed offshore penalty (which reflects the mean initial wealth) was around 3.5 times the median (GAO, 2013: 12), suggesting a heavy degree of skewness of initial wealth. In panels (e) and (f) we report the results of an exercise in which we increase the skewness of initial wealth by increasing the parameter α in (7) while adjusting β to hold aggregate offshore evasion constant. The effects of increasing skewness in this way affect. As skewness increases in this way, undisclosed liabilities becomes more concentrated among a smaller set of taxpayers, allowing the tax authority to focus its auditing effort on a smaller number of taxpayers (by becoming more lenient). To offset the negative effect on accepted disclosures of becoming more lenient, the tax authority simultaneously reduces the incentive for honesty.

We next consider the effects of the preferences and beliefs of an individual taxpayer, in particular the curvature of the utility function (y) and the degree of uncertainty over the value of a (σ_a). Panels (g) and (h) of Figure 1 the effects of increasing the curvature of utility (risk aversion). When taxpayers are close to being risk neutral, undisclosed liabilities are significant enough that the tax authority audits all disclosures and offers no incentive for honesty. As taxpayers become more risk averse, offshore evasion decreases, and disclosures increase, implying dwindling undisclosed liabilities. The tax authority responds both by becoming *more lenient* and by offering *stronger incentives* for honesty.

Unlike the taxpayer's utility function, which is wholly outside the control of the tax authority, the degree of uncertainty of a taxpayer over the value of a may at least partially be influenced by the tax authority. At one extreme, the tax authority could communicate the value of a in its letter to taxpayers, and at the other extreme it could purposefully promote confusion over its value by sending contradictory signals. Does a net-revenue maximizing tax authority promote or dispel uncertainty over a ? Panels (i) and (j) of Figure 1 investigate this question. Panel (j) shows that *net-revenue is increasing in taxpayer uncertainty over a* . When uncertainty over a is low, taxpayers are relatively able to judge how much under-disclosure the tax authority will tolerate. Hence, most disclosures are accepted.²⁶ The tax authority principally reacts to growing uncertainty over how lenient it will be by becoming *more lenient*.²⁷ It also offers *stronger incentives* for honesty, but only on a relatively modest scale.

4.3.2. Tax authority variables

We now turn to the variables that characterize the tax authority: its cost of audit (c_H) and the quality of its signal (σ_g). Panels (k) and (l) of Figure 1 show the results for the cost of audit. As is intuitive, when c_H is very low the tax authority optimally audits every disclosure, and when c_H is very high it accepts every disclosure. On an intermediate range, however, the tax authority's growing reluctance to conduct audits sees it become *more lenient*, and offer *stronger incentives* for honesty. The quality of the tax authority's signal matters, for a lower quality signal causes it not to audit taxpayers that it should audit, and to audit taxpayers that it should not audit. Consequently, the tax authority's growing unwillingness to conduct audits as its signal worsens sees it become *more lenient*, and offer *stronger incentives* for honesty (panels (m) and (n) of Figure 1). Note that, on the interval in which f_A is unconstrained, these effects are sufficiently strong that disclosures are seen to be increasing in σ_g , thereby reversing the sign of the effect we find in section 4.2 (where a and f_A are fixed).

4.3.3. Fiscal/legal variables

We view the marginal tax rate and the unincentivized fine rate as two parameters set by government, but which the tax authority nevertheless treats as fixed when designing a IOVDS. As discussed in Rablen (2014) and elsewhere, tax authorities typically do not set tax rates, while fine rates are prescribed in legislation, making them slow and costly to change. How does a net-revenue maximizing tax authority set f_A as f_H increases? The results presented in panels (o) and (p) of Figure 1 imply that this ratio is increasing in f_H , such that the tax authority offers *weaker incentives* for honesty. The intuition for this finding is that a higher unincentivized fine rate increases the returns to audit activity, making the tax authority less keen to incentivize taxpayers to make a disclosure that will be accepted. This intuition also accounts for the finding that the tax authority also chooses to be *less lenient* in its choice of a .

²⁶ In the extreme version of the model, with uncertainty over neither a nor q , every taxpayer would make the (accepted) disclosure $x = \max(0, E - [a / \theta])$.

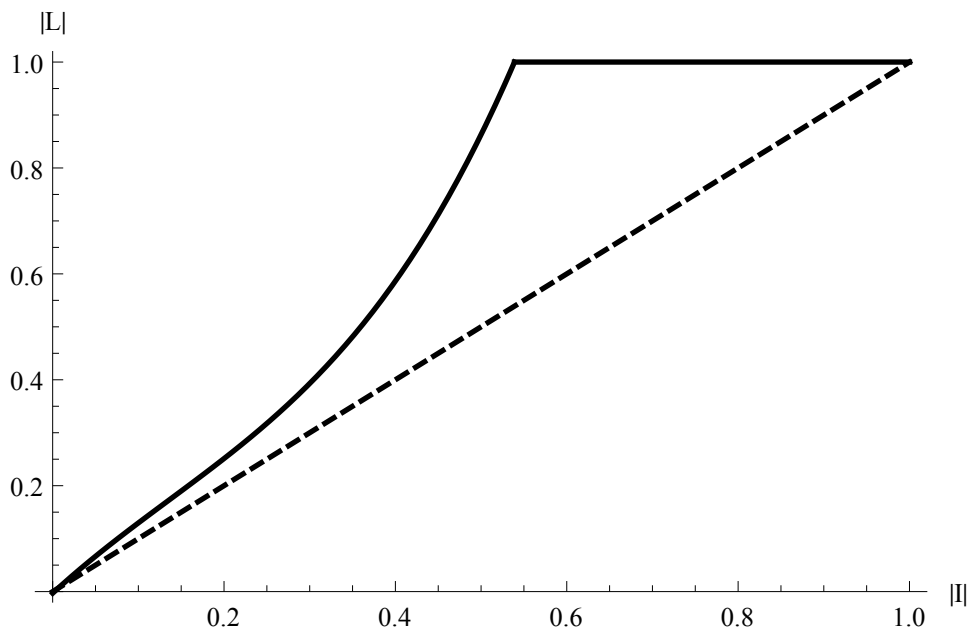
²⁷ Note in panel (j) of Figure 1 that, on the interval in which f_A is held constant at its maximum, the endogenous increase in a is sufficiently strong that disclosures are decreasing, overturning the positive effect found in section 4.2 for a fixed a .

As discussed in section 4.2, increases in the marginal tax rate drive up evasion. At very low levels of θ the tax authority performs no audits, as there is relatively little offshore evasion, and the tax liability associated with this evasion is also small. Conversely, at very high values of θ the tax authority audits all disclosures, for by this time there is significant offshore evasion carrying a substantial tax liability.

4.3.4. Choice of letter set

The final design aspect that we are able to address within the model is the choice of size of the letter set, conditional upon the size of the information set. We model this choice through the parameter κ , where $\kappa = 0$ implies that communication is only with taxpayers belonging to \mathbf{I} , and $\kappa = 1$ implies that communication is with all taxpayers in \mathbf{T} . We find (Figure 2) that the tax authority always finds it optimal to communicate with more taxpayers than it truly has information on, and the extent to which this is so is increasing in the size of \mathbf{I} . Thus, for $|\mathbf{I}|$ above around one-half, it is optimal for the tax authority to communicate with all taxpayers in \mathbf{T} . We therefore find some theoretical support for the approach of the IRS in suppressing the criteria used to determine which UBS customers had their offshore assets revealed. When \mathbf{I} is small relative to \mathbf{T} , however, the model suggests against communication with all taxpayers, as the revenue from additional disclosures from taxpayers outside \mathbf{I} is outweighed by the reduction in revenue from disclosures by taxpayers inside \mathbf{I} .

FIGURE 2. Optimal size of the letter set conditional on the size of the information set.



5. Conclusion

In this paper we present a theoretical model of IOVDS. We characterize the optimal behavior of taxpayers in such Schemes, and use this characterization to investigate the optimal choice of the key design parameters. Our main findings are that tax authorities can raise revenue while reducing their costs by running IOVDS; that shrouding the audit rule for disclosures under an IOVDS raises net revenue; and that some broadening of the base of IOVDS beyond taxpayers for whom information has been observed is always optimal, though broadening the base too much can ultimately reduce revenue. We also show how the optimal degree of incentivization is related to characteristics of the tax authority and of the taxpayers.

Although this study represents an important first step in our understanding of IOVDS, we believe that it represents no more than a preliminary exploration. In particular, the model relies heavily on the notion of pre-commitment, although it is unclear that tax authorities can commit in the way supposed. Three important next steps will be to: (1) endogenize the formation of beliefs, which in this model are exogenously asserted; (2) give taxpayers an explicit choice

as to whether or not they make a disclosure under an IOVDS (in the present framework the closest a taxpayer can get to not participating in a Scheme is to disclose zero); and (3) analyze the choice of letter set without recourse to the assumption that disclosures from taxpayers outside the information set are necessarily accepted. We hope to take up these challenges in our future research.

Bibliography

- Allingham, M.G., and Sandmo, A. (1972). Income tax evasion: A theoretical analysis, *Journal of Public Economics* 1 (3–4): 323–338.
- Alm, J., and Beck, W. (1993). Tax amnesties and compliance in the long run: A time series analysis, *National Tax Journal* 46(1): 53–60.
- Alm, J., McKee, M., and Beck, W. (1990). Amazing grace: Tax amnesties and compliance, *National Tax Journal* 43(1): 23–37.
- Andreoni, J. (1991). The desirability of a permanent tax amnesty, *Journal of Public Economics* 45(2): 143–159.
- Arrow, K.J. (1971). *Essays in the Theory of Risk-bearing*, Chicago: Markham Publishing.
- Bayer, R.-C. (2006). Moral constraints and evasion of income tax, *IUP Journal of Public Finance* 4(1): 7–31.
- Becker, G.S. (1968). Crime and punishment: an economic approach, *Journal of Political Economy* 76(2): 169–217.
- Boesler, M. (2012). *The Controversial 'Lagarde List' Has Leaked, and It's Bad News for the Greek Prime Minister*, Business Insider, New York: Business Insider Inc. Available: <http://www.businessinsider.com>.
- Center for Public Integrity (2013). *Secrecy for Sale: Inside The Global Offshore Money Maze*, Washington: Center for Public Integrity.
- Coelho, R., Richmond, P., Barry, J., and Hutzler, S. (2008). Double power laws in income and wealth distributions, *Physica A* 387(15): 3847–3851.
- Committee of Public Accounts (2008). *HMRC: Tackling the Hidden Economy*, HC 712, London: The Stationery Office.
- Committee of Public Accounts (2012). *Public Accounts Committee – Minutes of Evidence*, HC 716, 5th November, London: The Stationery Office.
- European Union (2003). Council Directive 2003/48/EC of 3 June 2003 on taxation of savings income in the form of interest payments, *Official Journal of the European Union* L157: 38–48.
- Franzoni, L.A. (2000). Amnesties, settlements and optimal tax enforcement, *Economica* 67: 153–176.
- GAO (2013). *Offshore Tax Evasion: IRS has Collected Billions of Dollars, but May be Missing Continued Evasion*, GAO-13-318, Washington: Government Accountability Office.
- HMRC (2006). *HMRC Annual Report 2005-6*, London: HMRC.
- HMRC (2009). *New Disclosure Opportunity: Making a Disclosure*, NDO1, London: HMRC.
- HMRC (2012). *Penalties for Inaccuracies in Returns and Documents*, CC/FS7, London: HMRC.
- IRS (2013). *Whistleblower Office At-a-glance*, Washington, DC: IRS. Accessed at www.irs.gov/uac/Whistleblower-Office-At-a-Glance.
- Klass, O.S., Biham, O., Levy, M., Malcai, O., and Solomon, S. (2006). The Forbes 400 and the Pareto wealth distribution, *Economics Letters* 90(2): 290–295.
- Kolda, T.G., Lewis, R.M., and Torczon, V. (2003). Optimization by direct search: New perspectives on some classical and modern methods, *SIAM Review* 45(3): 385–482.
- Lee, K. (2001). Tax evasion and self-insurance, *Journal of Public Economics* 81(1), 73–81.
- Lewis, R.M., Shepherd, A., and Torczon, V. (2007). Implementing generating set search methods for linearly constrained minimization, *SIAM Journal on Scientific Computing* 29(6): 2507–2530.
- Macho-Stadler, I., and Pérez-Castrillo, J.D. (2002). Auditing with signals, *Economica* 69 (273): 1–20.
- Malik, A.S. and Schwab, R.M. (1991). The economics of tax amnesties, *Journal of Public Economics* 46(1): 29–49.

- Mealem, Y., Tobol, Y., and Yaniv, G. (2010). Whistle-blowers as a deterrent to tax evasion, *Public Finance Review* 38(3): 306–320.
- Melamud, N.D., and Mookherjee, D. (1989). Delegation as commitment: the case of income tax audits, *Rand Journal of Economics* 20(2): 139–163.
- Meyer, D.J., and Meyer, J. (2005). Relative risk aversion: what do we know?, *Journal of Risk and Uncertainty* 31(3): 243–262.
- Mookherjee, D., and Png, I.P.L. (1989). Optimal auditing, insurance and redistribution, *Quarterly Journal of Economics* 104(2): 399–415.
- Morton, S. (1993). Auditing for Fraud, *Accounting Review* 68(4): 825–839.
- Oates, J. (2008). *HMRC Pays Criminal for 'Tax Dodger' Discs*, The Register, London: Situation Publishing. Available: <http://www.theregister.co.uk>.
- OECD (2010). *Offshore Voluntary Disclosure: Comparative Analysis, Guidance and Policy Advice*, Paris: Organisation for Economic Co-operation and Development.
- Peev, G. (2012). *Tax Evaders Won't Be Prosecuted...and Will Stay Anonymous: HMRC Refuses to Name 500 Britons on 'Lagarde List'*, 2nd November, MailOnline, London: Associated Newspapers.
- Rablen, M.D. (2014). Audit probability versus effectiveness: The Beckerian approach revisited, *Journal of Public Economic Theory* 16(2): 322–342.
- Reinganum, J., and Wilde, L. (1985). Income tax compliance in a principal-agent framework, *Journal of Public Economics* 26(1): 1–18.
- Reinganum, J., and Wilde, L. (1986). Equilibrium verification and reporting policies in a model of tax compliance, *International Economic Review* 27(3): 739–760.
- Sanchez, I., and Sobel, J. (1993). Hierarchical design and enforcement of income tax policies, *Journal of Public Economics* 50(3): 345–369.
- Scotchmer, S. (1987). Audit classes and tax enforcement policy, *American Economic Review* 77(2): 229–233.
- Smith, H. (2013). *Greek Court Acquits Editor Who Leaked 'Lagarde List' of Suspected Tax Evaders*, The Guardian, 27th November, London: Guardian Media Group.
- Stella, P. (1991). An economic analysis of tax amnesties, *Journal of Public Economics* 46(3): 383–400.
- Thomson Reuters (2009). *ATO Extends Its Offshore Income Voluntary Disclosure Amnesty*, Thomson Reuters Weekly Tax Bulletin, 4th December, Thomson Reuters: New York.
- Treasury Committee (2012). *Closing the Tax Gap: HMRC's Record at Ensuring Tax Compliance*, HC 1371, London: The Stationery Office.
- Wakker P.P. (2008). Explaining the characteristics of the power (CRRA) utility family, *Health Economics* 17(12): 1329–1344.
- Watt, H., Winnett, R., and Newell, C. (2012). *HSBC Investigation: Clients of Britain's Biggest Bank Exposed*, The Telegraph, 10th November, London: Telegraph Media Group. Available: <http://www.telegraph.co.uk>.
- Yaniv, G. (2001). Revenge, tax informing, and the optimal bounty, *Journal of Public Economic Theory* 3(2): 225–233.

Appendix

Proof of Proposition 1 Assuming $E \in (0, w)$, we may totally differentiate the taxpayers' first order condition (foc1). Letting $\gamma = -xU''(x)/U'(x)$ be the coefficient of relative risk aversion, $v_0 = [1-\theta][w-E]$, and $v_1 = v_0 + E[1+r]$, we obtain

$$\begin{aligned}\frac{\partial E}{\partial w} &= \frac{E}{w} \frac{E[v + \gamma w[1-\theta]] + w[2v - \gamma w[1-\theta]]}{[1-\theta]}, \\ \frac{\partial E}{\partial \theta} &= \frac{E}{\theta} \frac{[w-E][v + \gamma \theta[w-E]]}{v[E+2w] + \theta \gamma E[w-E]} > 0; \\ \frac{\partial E}{\partial \beta} &= -\frac{E}{\beta} \frac{v[w-E]}{v[E+2w] + \theta \gamma E[w-E]} < 0; \\ \frac{\partial E}{\partial \gamma} &= -\frac{Ev[w-E]\log(v)}{v[E+2w] + \theta \gamma E[w-E]}.\end{aligned}$$

Hence

$$\begin{aligned}\frac{\partial E}{\partial w} > 0 &\Leftrightarrow \gamma < 2 + \frac{E}{w} \left[\frac{\theta E + w[3-\theta]}{[w-E][1-\theta]} \right] > 2; \\ \frac{\partial E}{\partial \gamma} < 0 &\Leftrightarrow v > 1.\end{aligned}$$

Therefore, for $\gamma \in [0, 2]$ we have $\partial E / \partial w > 0$. Noting that $v > w[1-\theta]$, if $\underline{w} \geq [1-\theta]^{-1}$ then $v > 1$ for all $w \geq \underline{w}$. We choose $\underline{w} > [1-\theta]^{-1}$, where $\bar{\theta} = 0.5$ is the highest value of θ considered in the analysis, hence $\partial E / \partial \gamma < 0$.