

The Use of Agent-Based Modelling to Investigate Tax Compliance

Nigar Hashimzade

Durham University and Institute for Fiscal Studies

Gareth D. Myles*

University of Exeter and Institute for Fiscal Studies

Frank Page

Indiana University

Matthew D. Rablen

Brunel University and Institute for Fiscal Studies

November 2013

Abstract

Agent-based modelling can be used to investigate the behavioural and social aspects of tax compliance. We illustrate the approach with two models. The first model emphasises the role of occupational choice in tax compliance, and explores the effect of non-compliance on risk-taking and income distribution. The modelling of the compliance decision is discussed with an emphasis on decision-making under uncertainty and social interaction. We then add to the model a social network which governs the transmission of information on attitudes and beliefs, and investigate alternative audit strategies. A strategy of auditing a fixed number of taxpayers from each occupation dominates alternative strategies (including random and focussed strategies) in the sense of first-order stochastic dominance.

Keywords: agent-based modelling, tax evasion, attitudes, beliefs, social network, occupational choice.

JEL classification: H26, D85, C63.

*Contact details: Gareth D. Myles, Department of Economics, University of Exeter, Exeter EX4 4PU, UK, email: gdmyles@ex.ac.uk. Thanks are due to the ESRC for financial support under grant ES/K005944/1. Previous versions of the paper were presented at the Ottawa Workshop on Compliance and at shadow2013 in Münster. The paper was completed during a visit to Bogazici University; their hospitality was appreciated.

1 Introduction

The economic analysis of tax compliance has the objectives of explaining and predicting compliance behaviour. Achievement of these objectives is essential for the design of beneficial interventions that increase the level of compliance and raise revenue. There are several different research methodologies that can contribute to this programme of research. Theoretical analysis can develop models that are evaluated by empirical studies and tested using lab and field experiments. The focus of this paper is a further methodology that can be usefully applied to analyze compliance: *agent-based modelling*. This is a research methodology that is steadily gaining in popularity due to its flexibility and potential sophistication. We hope that the paper will demonstrate that agent-based modelling can yield fresh insights when applied to the study of compliance.

A successful application of agent-based modelling uses the best of economic theory to describe the behaviour of agents with heterogeneous characteristics and allows for interaction among these agents in a rich environment. The components of economic theory on which we focus are recent behavioural advances in understanding of the compliance decision, the effect of occupational choice in creating opportunities for non-compliance, and the role of social networks in the transmission of information. In brief, our model of the compliance decision and policy intervention combines *attitudes* towards compliance, *beliefs* about audit strategy, and *opportunities* for evasion. It also recognizes the *social setting* in which the compliance decision is made.

The paper describes the theoretical background of the modelling and the numerical results from two different agent-based models. The first model focuses on occupational choice and the distributional consequences of non-compliance. The second model generalizes the first by adding repeated social interaction and the transmission of attitudes and beliefs in a dynamic setting. The models demonstrate that non-compliance increases

inequality and risk-taking in the economy, and that different compliance behaviours can be established within occupational groups. It is also possible that taxpayers, on average, can systematically hold a belief about the probability of audit that remains consistently above the true rate. When audit strategies are compared we find that a strategy of auditing a fixed number of individuals within each occupational group delivers a higher level of revenue than strategies with randomness across groups or a systematic focus on groups.

Section 2 provides a descriptive introduction to agent-based modelling. Successful application of agent-based modelling requires a credible model of individual choice. In our context the role of *opportunities* for non-compliance is central. Section 3 consequently implements an agent-based model with choice of occupation using an extension of the Allingham and Sandmo (1972) framework. We then extend the model further to incorporate advances from behavioural economics, including the endogenous development of *attitudes* and *beliefs* within a *social network*. Section 4 reviews the literature on behavioural explanations of the individual compliance decision, and Section 5 describes how the behavioural concepts are implemented in the model. The paper is completed in Section 6 by analyzing the choice of *audit strategy* in an agent-based model that includes network effects and behavioural assumptions on preferences. Section 7 concludes the paper.

2 Agent-Based Modelling

Agent-based modelling is a computer simulation technique that is increasing in popularity for the study of economic and social behaviour. It involves the construction of a set of agents and an environment in which they interact, and has proved useful in many different areas of natural science and social science. There have been numerous economic applications (surveyed in Tesfatsion, 2006) and several previous studies of tax compliance (Andrei et al., in press; Bloomquist, 2004, Davis et al., 2003, Hashimzade et al., in press;

Korobow et al., 2007). Before describing what our work contributes to this literature, we provide in this section a general introduction to agent-based modelling.

To implement an agent-based model the first step is to define the agents that will interact and the environment in which the interaction takes place. In general, an individual agent will be characterized by their ability, objective, and information set. Some of the characteristics will be fixed at the outset of the simulation (e.g. ability) but others may be updated by experience (e.g. information). In economic applications agents are typically assigned an objective such as maximization of income or utility and make choices to achieve the objective. This need not be the case, and in many other areas of science agents can be mechanistic (e.g. interaction of particles controlled by the laws of dynamics) or simply random (e.g. very basic biological interaction). The number of agents and the distribution of characteristics of agents can be chosen according to the context of the research question or selected by a random process.

The second step is to construct the environment within which the agents interact. An economic environment could be a market place with trading rules or an economy with some set of institutions that govern interaction. An application in physics may involve placing particles in a dust cloud, or placing animals in a field for an application in biology. There may also be randomness involved in the choice of the environment.

Given the agents and the environment, the final step is to allow the agents to interact (economic agents can buy and sell, or particles collide and coalesce) and to observe the outcome. If there are multiple periods of interaction then both the dynamic process and its steady states can be of interest. These will be governed by the initial state of the system, by the choices made by the agents, and by any random components during the interaction. The parameters of the system, or the probability distributions governing choice of parameters, can be varied to test the effect of their choice on the outcome. The government can either be an agent that chooses policy, or else it can be part of the

environment with policy as a parameter.

A classic agent-based model of nature is that of the predator-prey relationship. An implementation of this is available in the free Netlogo software (<http://ccl.northwestern.edu/netlogo/>) that provides a platform for developing agent-based models. The Wolf Sheep Predation model is set in a grassy landscape. The sheep wander randomly around the landscape eating grass and reproducing. The wolves also move randomly around the landscape until they encounter sheep which they devour. Movement costs wolves energy so they must eat sheep to survive. If a wolf does not find any sheep it will eventually run out of energy and die. The purpose of the simulation is to describe the evolution of the two populations.

The Wolf Sheep Predation model is helpful for raising the question of what is learnt from a simulation. In this case there is no need to run the model to predict what the long-run outcome must be since there are only two steady-states of the model. One steady state involves a world populated only by sheep who can survive on grass in the absence of wolves. The other steady state involves no life whatsoever. Which steady state arises is dependent on random events within the simulation. The initial conditions and the parameter values determine the likelihood of each steady state. In both cases the wolves become extinct; what differs is whether the wolves eat all the sheep or not before facing extinction. The simulation does not add anything *qualitative* to this knowledge of the steady states. Whether it can add anything *quantitative* (such as the time taken for wolves to become extinct) is very debatable and depends on the validity of the calibration. For the simple model described it is highly unlikely that there can be anything of quantitative significance. What these comments illustrate is that a simulation is not an end in itself, but is only justified if it provides insight that could not be obtained by any alternative analytical means.

A range of free software is available for undertaking agent-based simulations. For an economist the usefulness of this software is typically limited by the fact that it does

not permit agents to undertake complex optimization within the simulations. This is important in many applications since it is the inclusion of optimal choice that distinguishes economic behaviour from modelling in the natural sciences. For some models it may be possible to compute explicit solutions to the optimization problems in which case the free software is adequate. Whenever a numerical optimization sub-routine is required, as it is in the models described in the following sections, it is necessary to employ suitable software (such as Matlab) for writing dedicated codes.

3 Risk-Taking and Income Distribution

A key element for understanding the compliance decision is the role the *opportunity* for non-compliance plays in the choice of occupation. Working as a paid employee either rules out non-compliance, if labour income is subject to a withholding tax (such as the PAYE system in the UK), or makes successful non-compliance very unlikely, if there is a system of third-party reporting. In contrast, choosing to be self-employed and accepting the responsibility for tax filing opens the opportunity for non-compliance. It is through this channel that occupational choice is inter-linked with the compliance decision.

There is a second aspect of occupational choice that is also linked to the compliance decision. It is generally true that the level of income received from employment is more certain than the income that will be generated from self-employment. This implies that choosing self-employment also involves accepting greater income risk and, therefore, all else constant, the self-employed will have a lower degree of risk aversion than the employed. This directly determines the extent of non-compliance: the amount of income that is not declared increases as risk aversion decreases. In this way occupational choice self-selects those who will evade most into an occupation where they have the opportunity to evade.

Our first example of agent-based modelling incorporates occupational choice into a

compliance model. This is achieved by extending the model of Allingham and Sandmo (1972) to permit each individual, first, to make an occupational choice and, second, to make an evasion decision based on the realization of income. The model can be seen as a generalization of the work of Pestieau and Possen (1991). The focus of the simulation is the effect that non-compliance has upon the amount of risk-taking in the economy and income distribution.

The model has three occupations. *Employment* is modelled as a safe occupation with a fixed wage that can differ among individuals. There is no opportunity to be non-compliant in employment due either to the operation of a withholding tax or through third-party reporting. The other two occupations are different forms of *self-employment*. Self-employment is intended to represent running a small business, and so the income is assumed to be risky. However, it is possible to evade tax on income from self-employment since it is not subject to the same degree of third-party reporting. We adopt the natural assumption that each self-employed person makes a compliance decision after the (random) income from self-employment is realized. The choice of occupation is made by comparing the utility derived from employment to the expected utility (taking into account optimal compliance for each income realization) from the two self-employment occupations. The occupation that delivers the highest utility is chosen. In this simulation the occupational choice decision is made once. In the simulations of sections 5 and 6 it is made at the start of every period because the choice may change as the information of the taxpayer evolves through the interaction with others.

The simulation randomly assigns to each taxpayer a set of four characteristics, $\{w, \rho, s_1, s_2\}$, where w is the wage in employment, ρ is the coefficient of (relative) risk aversion in a constant relative risk aversion (CRRA) utility function, and s_i is the level of skill in self-employment occupation i . The income earned from self-employment in occupation i is $s_i y_i$

where y_i is drawn from a beta distribution $g(\cdot)$.¹ The variable y_i can be interpreted as local market conditions, so that income is determined jointly by individual skill and market conditions. The draw of y_i is unique for each taxpayer, so in a given round of simulation, a low-skill individual in occupation i may earn more than a high-skill individual if the former obtains a beneficial draw of y_i . It is assumed that $\mu(y_1) < \mu(y_2)$ and $\sigma^2(y_1) < \sigma^2(y_2)$, so that for a given skill level self-employed occupation 2 has a higher mean income but also a greater variance of income. We therefore refer to occupation 2 as being riskier than occupation 1. If a taxpayer has realized outcome $s_i y_i$ from self-employment i the amount of income that is not declared, $E_i(y_i)$, is determined by

$$\max_{\{E_i\}} U(E_i; y_i) = pU([1-t]s_i y_i - ftE_i) + (1-p)U([1-t]s_i y_i + tE_i).$$

Taking account of the choice of E_i , the expected utility from self-employed occupation i is then

$$\mathcal{E}U_i = \int U(E_i(y_i)_i; y_i) g(y_i) dy_i.$$

The expected payoffs from the three occupations $\{U_0, \mathcal{E}U_1, \mathcal{E}U_2\}$ are compared (where occupation 0 is employment), and the maximum payoff determines the chosen occupation.

The agent-based simulation performs the following steps:

1. Individual characteristics are randomly drawn;
2. Occupation is chosen given characteristics;
3. Incomes are realized and the compliance decision is made;
4. The tax authority conducts random audits and punishes any evasion that is detected.

¹The advantage of the beta distribution is in the flexible choice of parameters, allowing density functions with required skewness to be obtained, and the finite support ensures robust convergence of the numerical integration.

The outcome is calculated for two different scenarios. The first scenario assumes that all income is honestly declared. This provides a baseline from which to judge the effect of non-compliance. The second scenario assumes that non-compliance takes place. Each simulation has 1000 individuals and is repeated 100 times. The data are averaged across the 100 rounds in order to smooth out the consequences of randomness. The following parameters were used for an illustrative example: tax rate is 25 percent, each self-employed taxpayer is audited with probability of 5 percent, and the fine rate is 150 percent of evaded tax.

Our first two figures compare the distribution of occupational choices between the two scenarios. Figure 1 shows the distribution of taxpayers across the three occupations with honesty. The three occupations are on the horizontal axis and the vertical axis shows the number of taxpayers in each occupation. The distribution when non-compliance is possible is shown in Figure 2. Comparing the figures shows that non-compliance causes the distribution of occupational choices to shift away from employment toward the two risky self-employment occupations. As a consequence there is more occupational risk-taking when non-compliance is possible. In addition to this increase in occupational risk-taking there is a further increase in total risk-taking in the economy because some of the taxpayers choosing self-employment are also evading. Hence, the total amount of risk-taking in the economy is increased by the existence of tax evasion. This observation is interesting in view of past discussion (Kanbur 1981, Black and de Meza 1997) on the efficiency of risk-taking in competitive economies.

The effect of non-compliance on income distribution is presented in two ways. Table 1 provides summary statistics of the income distributions with and without evasion, and figure 3 plots the Lorenz curves for the two distributions. The effect of non-compliance is to increase the mean income level, where the mean is computed after both taxes and fines have been imposed. Non-compliance also increases the inequality of income as measured




Figure 1: Occupational choice with honest tax payment




Figure 2: Occupational choice with non-compliance

	Honesty	Non-compliance
Mean income	9.986	13.671
Gini coefficient	0.380	0.428

Table 1: Income distribution




Figure 3: Lorenz Curves for honesty (solid) and non-compliance (dashed)

by the Gini coefficient. Figure 3 shows that there is Lorenz-curve dominance for the income distribution with honesty, and so the ranking is independent of the inequality index.

Another consequence of non-compliance is that those who fail to declare their true income do not pay the statutory tax rate. Define the effective tax rate for a non-compliant taxpayer who is not audited by

$$ETR^{NA} = \frac{\text{Tax payment on income declared}}{\text{Actual income}}, \quad (1)$$

and for a non-compliant taxpayer who is audited by

$$ETR^A = \frac{\text{Correct tax payment plus fine}}{\text{Actual income}}. \quad (2)$$

ETR^{NA} will be below the statutory tax rate and ETR^A will be above the statutory tax rate. The consequence of non-compliance by taxpayers is that the distribution of effective tax rates is unrelated to income and does not correspond to the flat tax intended by the government. This point is illustrated in figure 4 which displays a histogram of tax rates. This is tri-modal, reflecting the three groups: non-compliant taxpayers who are




Figure 4: Distribution of tax rates

not audited, compliant taxpayers, and non-compliant taxpayers who are audited. Given the propensity for taxpayers to be non-compliant and the audit rate of 5 percent, the majority of taxpayers pay an effective tax rate below the statutory rate of 25 percent. The general observation is that non-compliance undermines the intended tax policy of the government.

These results illustrate some of the effects that non-compliance can have upon the economy. The possibility of non-compliance encourages entry into risky occupations, while the consequence of non-compliance and auditing is increased inequality and a dispersion of the effective tax rate. The agent-based model reported in this section shows the importance of introducing opportunities, but there are more features of the compliance decision that need to be taken into account. The next section therefore reviews recent literature on the applications of behavioural economics to the compliance decision.

4 Modelling Compliance

The properties of an agent-based model are determined by the behaviour of the individual agents. This implies that modelling the *choice behaviour* behind the compliance decision is key to obtaining interesting and credible insights. The aim when constructing a model

should be to integrate the best of current theory and evidence. For the compliance decision this involves an acknowledgement of the limitations of the Allingham and Sandmo (1972) model of tax compliance and the incorporation of ideas from behavioural economics. The purpose of this section is to briefly review some models of the individual compliance decision. A more complete survey can be found in Hashimzade et al. (2013).

Research on compliance behaviour has built on the Yitzhaki (1974) model which was itself a modification of Allingham and Sandmo (1972). Correspondingly, we refer to this below as the *ASY model*. The amount of evasion, E , is chosen to maximize expected utility

$$\mathcal{E}U = pU(Y[1-t] - t f E) + [1-p]U(Y[1-t] + t E),$$

where p is the probability of audit, Y is income, t is the tax rate, and f is the fine levied on tax evaded. The model takes the level of income as fixed. As we have already noted, the source of income is an important determinant of the *opportunity* for evasion due to third-party reporting and withholding on employment income. The inclusion of occupational choice is one of the central features of our agent-based models. For the present, we set this issue aside and focus on choices contingent on income.

The literature has identified two problems with the predictions of the ASY model. First, when confronted with the parameter values observed in practice the model predicts that *all* taxpayers should be non-compliant. Formally, the necessary and sufficient condition for $E > 0$ is

$$p < \frac{1}{1+f}.$$

In practice, the value of f is rarely more than 2, so non-compliance occurs ($E > 0$) if $p < 1/3$. The exact value of p is information that only revenue services are party to, but no revenue service audits anywhere are even approaching one third of taxpayers. In this sense, all taxpayers should be non-compliant. Second, the predicted relationship between

the amount of non-compliance and the tax rate is counter to intuitive expectation and counter to some (but not all) evidence. The formal result is that decreasing absolute risk aversion is a sufficient condition for

$$\frac{dE}{dt} < 0.$$

These results have led to a considerable research effort to identify alternative models of the compliance decision that make predictions with greater conformity to the facts. The solutions proposed to improve the predictions of the model include appeal to *non-expected utility theory* and to *social customs*. We will discuss each of these in turn.

A general representation of non-expected utility choice theory is given by writing the value function, V , as

$$V = w_1(p, 1 - p)v(Y[1 - t] - t f E) + w_2(p, 1 - p)v(Y[1 - t] + t E). \quad (3)$$

In (3) $w_i(p, 1 - p)$, $i = 1, 2$, are *weighting functions* that translate the probabilities p and $1 - p$ into more general weights. The typical assumption is that unlikely events are overweighted, so in the context of compliance $w_1(p, 1 - p) > p$. The function $v(\cdot)$ is a payoff function that can be more general than a utility function. For example, it is normally assumed that utility is concave ($U'' < 0$) which is not a property that a value function need satisfy.

Within this general framework several alternatives have been proposed:

- *Rank Dependent Expected Utility* (Quiggin 1982) imposes structure on the weighting functions
- *Prospect Theory* (Kahneman and Tversky 1979) uses weighted probabilities and an *s-shaped* payoff function, and compares incomes to a reference point
- *Non-Additive Probabilities* (e.g., Chateauneuf 1994) do not require the normal consistency of aggregation for probabilities

- *Ambiguity* (e.g., Snow and Warren 2005) permits uncertainty over the probability of outcomes

The appearance of weighting functions (or *beliefs*) in these alternative preference structure can improve the predictions by making the sufficient condition for evasion tighter and individual-specific. However, they do not change the direction of the tax effect to make $dE/dt > 0$. In addition, these alternatives can have their own shortcomings as explored in detail in Hashimzade et al. (2013). Variants of prospect theory to describe tax compliance are used, for example, by Yaniv (1999), al Nowaihi and Dhami (2007), and Bernasconi and Zanardi (2004). A difficulty with this approach can be seen by adopting the standard Kahneman-Tversky value function

$$v(z) = \begin{cases} z^\beta, & \text{if } z \geq 0, \\ -\gamma(-z^\beta), & \gamma > 1, \text{ if } z < 0, \end{cases} \quad (4)$$

and choosing the reference point as income if the correct tax payment is made, $Y[1 - t]$. The payoff function then becomes

$$V = E^\beta t^\beta [w_2 - w_1 \gamma f^\beta], \quad (5)$$

so that the optimal choice is either to comply in full or to declare no income. This is a simple consequence of the non-concavity of the objective function.

The existence of stigma from non-compliance and the existence of a social custom for compliance have identical formal representations. Correspondingly, we focus on social customs in what follows. A social custom is an informal rule of behaviour that summarizes the *attitude* toward compliance. A loss of social custom utility (or alternatively, a stigma cost or psychic cost) is incurred if the custom is broken

$$V = \begin{cases} U(Y[1 - t]), & \text{if } E = 0, \\ \mathcal{E}U - \chi^i, & \text{if } E > 0. \end{cases} \quad (6)$$

Across individuals there will be a cutoff χ^* such that $\chi^i < \chi^* \implies E > 0$ and $\chi^i > \chi^* \implies E = 0$. If $\chi^i = \chi^i(m, E)$, (m the proportion of population evading) evasion becomes a

social decision. Myles and Naylor (1996) show that $\chi_m^i(m, E) < 0$ opens the possibility of multiple equilibria.

For some specifications of the stigma or social cost, it becomes possible to obtain $dE/dt > 0$. In a recent paper Piolatto and Rablen (2013) disentangle four distinct elements of prospect theory in their roles for the individual compliance decision; in particular, they find that probability weighting has no effect upon the sign of dE/dt . Furthermore, they prove that when the expected utility theory model is augmented with stigma, or the psychic cost of non-compliance, it can overturn the sign of the tax effect. Thus, prospect theory offers no fundamental advantage over the expected utility theory with this particular modification. Based on these observations, we do not need to feel bound by using either expected utility using objective probability or to be restricted by any of the particular alternatives to the expected utility theory. We proceed, therefore, by mixing subjective beliefs and stigma with convenient functional forms.

5 Attitudes, Beliefs and Network Effects

The empirical analysis of the determinants of tax evasion has demonstrated two important features. First, there is a strong evidence that the social setting influences the individual compliance decision. For example, individual perceptions of the justifiability of tax evasion in a country are positively associated with the measures of aggregate tax evasion in that country, according to the World Values Survey (Slemrod 2007). We refer to this effect as the *attitude to compliance*, or *attitude*. An aggregate measure of the individual attitudes to compliance across a society can also be viewed as the tax morale prevailing in that society. One can think about the effect of tax morale upon the individual attitudes to compliance as an externality: an individual who holds the view that non-compliance can be (sometimes) justified contributes to the low tax morale in the society which, in turn,

makes for that individual the decision to evade tax more easily acceptable.

Second, the probability of audit is not revealed to taxpayers by the revenue service. Therefore, in the individual evaluation of the expected benefit from evasion the probability of being audited and found to be non-compliant is subjective, rather than objective. While the objective probability is part of the audit strategy of the revenue service, the subjective probabilities may be formed on the basis of individual experience and available information, and can, of course, be different for different individuals. To distinguish between the objective and the subjective probabilities we refer to the latter as the *subjective belief*, or just *belief*.

If attitudes and beliefs are determined, among other factors, by experience and information, it is natural to assume that they can evolve and change for a given individual over time as he or she interacts with the revenue service (through accumulation of experience) and with other individuals (through accumulation and exchange of information). Furthermore, information exchange is more likely to occur (or more information is likely to be exchanged) if the individuals belong to the same occupational group. Individuals meet with their contacts in the social network, and meetings allow exchange of information on beliefs. In addition, at a meeting, individuals may observe each other's attitude to evasion. For example, individual 1 can infer something about individual 2's attitude simply by learning whether or not individual 2 has evaded tax previously. This will affect 1's own attitude and, through this channel, 1's future evasion decisions. The same may take place for individual 2. This, in particular, can explain why social groups have different behaviour with respect to tax evasion.

We have incorporated the dynamics of attitudes and beliefs into an agent-based model by adding to the individual compliance decision a process of *learning within a social network* according to the algorithm outlined below.

5.1 Networks and meetings

In an economy with N individuals the social network is described by a symmetric $N \times N$ matrix A with $A_{ij} = 1$ if individuals i and j are linked and $A_{ij} = 0$ otherwise. The links are bi-directional: if i “knows” j then j “knows” i .² In our simulations the network is fixed at the outset and does not change; one can also introduce random or endogenous changes in the network structure. Time is divided into discrete periods, and in every period each individual chooses an occupation, earns income, and decides how much of this income to declare. Declarations are audited (according to some randomizing device as described below), after which individuals linked in the network randomly meet and exchange information.

Here we introduce two additional layers of randomness: not all individuals in the network meet in every period, and not every meeting results in an information exchange. This is implemented by introducing an $N \times N$ matrix C of zeros and ones, drawn randomly in each period; this matrix represents the probabilities of meetings between individuals. Thus, in each period a random selection of meetings occur described by an element-by-element product of A and C : individuals i and j meet during a period if $A_{ij}C_{ij} = 1$ and do not meet otherwise.

Furthermore, at a meeting of i and j information is exchanged only with some probability. It is possible to consider various patterns in the probability of information exchange; one plausible assumption is that the probability depends on the occupational groups to which i and j belong. More specifically, we assume that the probability of information exchange between i and j is higher when i and j belong to the same occupational group, and that it does not depend on their individual characteristics or other model parameters. With three occupations, in general, six different probabilities can be introduced, denoted

²A matrix that is not symmetric captures uni-directional links. This can be used to investigate the effect of a “celebrity”.

by $q_{\alpha\beta}$, where $\alpha, \beta \in \{e, 1, 2\}$, and $q_{\alpha\alpha} > q_{\alpha\beta}$ for all α and $\beta \neq \alpha$.

5.2 Formation of beliefs

In period t individual i makes an occupational choice and (after income is realized) a compliance decision on the basis of the subjective belief, p_t^i , that i will be audited and caught if non-compliant. The belief is determined by audits prior to t (experience) and interaction with other individuals (information).

5.2.1 Audits and beliefs

The first updating effect of an audit is that of experience and is described by

$$\tilde{p}_t^i = X_t^i P + (1 - X_t^i) d(p_t^i), \quad P \in [0, 1],$$

where $X_t^i = 1$ if i was audited in t and $X_t^i = 0$ otherwise. The choice of P and $d(p_t^i)$ allows flexibility in the modelling of the updating rule.

Two different mechanisms for the formation of subjective beliefs are considered, the *target effect* and the *bomb-crater effect*. With the target effect, immediately after an audit the subjective belief rises, possibly to one, and then decays. In other words, if i is audited and caught in period t he believes that now the tax authority will target him as an evader and will certainly or nearly certainly audit him again, but if i is not audited he believes that he is less likely to be a target and is less likely to be audited next time. In the simulations we assume the maximal target effect and proportional decay: $P = 1$ and $d(p_t^i) = \delta p_t^i$, $\delta \in (0, 1)$. With the bomb-crater effect (e.g., Guala and Mittone, 2005), immediately after an audit the belief falls, possibly to zero, and then rises. That is, if i is audited and caught in period t he believes that he is less likely to be audited again (a bomb is unlikely to hit a crater made by the previous bomb), but subsequently worries that his turn to be audited again is approaching. In the simulations we assume the maximal bomb-crater effect and proportional increase: $P = 0$ and $d(p_t^i) = p_t^i + \delta(1 - p_t^i)$,

$\delta \in (0, 1)$. The empirical evidence on which mechanism is correct is mixed and does not provide a decisive argument in favour of one over the other.

5.2.2 Information exchange and beliefs

The second updating effect takes place at a meeting. Specifically, the individuals meet after audits take place, and their own subjective belief is updated, either according to the target mechanism or to the bomb-crater mechanism. If an information exchange occurs at a meeting between i and j , i 's belief is further updated according to the rule

$$p_{t+1}^i = \mu \tilde{p}_t^i + (1 - \mu) [X_t^j P + (1 - X_t^j) \tilde{p}_t^j].$$

This can also be written

$$p_{t+1}^i = \begin{cases} \mu \tilde{p}_t^i + (1 - \mu) P, & \text{if } j \text{ audited at } t, \\ \mu \tilde{p}_t^i + (1 - \mu) \tilde{p}_t^j, & \text{otherwise.} \end{cases}$$

The belief p_{t+1}^i is carried into the next period and is used when making occupational choice and evasion decision. A similar update takes place for individual j . Beliefs at time $t = 0$ in the simulations are assigned randomly.

5.3 Formation of attitudes

In the social custom approach to individual decision-making it is assumed that an individual derives additional utility if his or her decision is in line with the social custom (equivalently, utility is lost if the decision goes against social custom). In general, the importance of the social custom, or its weight in the utility function, can be specific for an individual. Since a social custom emerges in a society of interacting individuals, it is reasonable to assume that the weight assigned to the social custom by an individual is determined by interaction in the social network. For example, if the social custom is to pay taxes honestly, the weight will be higher when the number of honest taxpayers known to that individual is greater.

In the simulation the dynamic process for the importance of the social custom is implemented in the following way. Each individual i is randomly assigned a level of importance, χ_0^i , at time $t = 0$. This value is then updated in those time periods when there is an information exchange between individual i and some other individual, say, j . The updating process is described by

$$\chi_{t+1}^i = \frac{1}{X(i) + 1} \left[\chi_t^i X(i) + \mathbf{1}_{[E_t^j=0]} \right],$$

where $X(i)$ is the number of previous meetings for i at which information was exchanged, and $\mathbf{1}_{[A]} = 1$ if A is true and zero otherwise. One can easily verify that in this formulation $\chi_{t+1}^i > \chi_t^i$ if information is exchanged with an honest taxpayer and $\chi_{t+1}^i < \chi_t^i$ if information is exchanged with a non-compliant taxpayer. This form of social custom is added to preferences over income described by a CRRA utility function.

5.4 Equilibrium

Having specified individual decision-making and the process of interaction with other individuals, we now turn to the audit strategy of the tax authority. As the benchmark case, we first assume a standard random probability of audit: each self-employed individual is audited with the same constant probability; those in paid employment are not audited.³

We ran simulations for an economy populated by $N = 1000$ agents with heterogeneous individual characteristics. Each agent is characterized by risk preferences (captured by the coefficient of relative risk aversion), wage in employment, skill level in the two self-employment occupations, a subjective probability of audit, and a weighting of the social custom. As with the simulation in section 3, risk preferences, wage in employment, and skills in self-employment are drawn at the outset and remain fixed for each agent.

³It is assumed tax authority knows that in paid employment income tax is fully deducted at source, and there is no opportunity for earning additional income that could be concealed. This assumption could be modified in a more general model to allow an additional income for individuals in employment and a possibility to evade tax on that income.




Figure 5: Risk aversion

Furthermore, we retain the assumption that earnings in self-employment are random and that self-employment occupation 2 has a higher mean and variance for equal skills levels. The subjective probability and the weight on the social custom are updated each period as described above.

At the beginning of every period an agent chooses an occupation, and, if self-employment is chosen, then observes an income realization and decides how much income to declare. Income declarations are randomly audited, and non-compliant taxpayers are fined if caught. Agents update their beliefs about audits, meet in the social networks and exchange information (with some probability). This information is used for the secondary update of beliefs as well as for adjusting the attitude to evasion. The process repeats in the next period. At time zero beliefs and attitudes are assigned at random; in the simulations the effect of the initial condition disappears after about 20 periods. The results of the simulations are reported for beta distribution of earnings.

Figure 5 illustrates the self-selection of individuals into different occupational groups according to their risk aversion, with time periods on the horizontal axis and risk aversion on the vertical axis. In the simulations each individual agent is assigned a coefficient of relative risk aversion drawn from a uniform distribution between 0 and 10. This remains




Figure 6: Subjective beliefs

constant for a given individual, and is one of the drivers of the occupational choice and the compliance decision. The red, green and blue lines show the average coefficient of risk aversion of agents in paid employment and self-employed occupations 1 and 2, respectively. The averages move around because agents switch between occupations as their beliefs about audits and attitudes to evasion change over time. There is a clear indication that agents with higher risk aversion choose paid employment, whereas the agents with the lowest risk aversion choose the riskier type of self-employment.

The patterns in beliefs and attitudes that emerge in the economy are illustrated in Figures 6 and 7. Both the belief about audits and the attitude to evasion are affected by the interaction and information exchange in social networks, and information exchange is more likely between those members of social network who are in the same occupation. It is expected, therefore, that agents in different occupations are likely to exhibit, on average, different beliefs and attitudes. Figure 6 shows that, indeed, the subjective belief about the probability of audit is the lowest among employed, close to the objective probability at 0.05 (in these simulations), whereas for self-employed it is sustained at a much higher level, about 0.18. This pattern obtains under the target effect assumption on the belief update; under the bomb-crater effect the subjective beliefs are persistent at an even higher




Figure 7: Occupations and compliance

level.

Figure 7 illustrates the average rate of compliance (the proportion of honest declarations) for each type of self-employment, along with the economy-wide rate of compliance. Compliance is lower in the riskier occupation: just over 30 percent of agents in self-employment 2 declare their income honestly, whereas in self-employment 1 the rate is around 50 percent; this illustrates our point about the link between risk-taking in the choice of occupation and in the evasion decision. The overall level of compliance in the economy is around 62 percent. The differences in compliance rates are driven partly by the differences in risk aversion and partly by the differences in attitudes to evasion: exchanging information with more compliant agents reinforces the importance of the social norm of compliance, and, conversely, interacting with non-compliant agents makes non-compliance feel less unacceptable.

6 Audit Strategies

The model is sufficiently rich to permit a range of questions to be investigated. Of particular interest is the choice of audit strategy by the tax authority. Audits are costly,

and the tax authority might be interested in identifying a strategy or a set of strategies that deliver the highest compliance at a given cost, or result in the highest revenue collected net of audit cost. It is natural to ask, for example, whether random audits or audits targeting a particular group of taxpayers, or some mix of both, deliver a higher tax yield in an environment where taxpayers are influenced in their compliance decisions by their own experience as well as the experience of other taxpayers. In addition to the benchmark case of random audits with constant exogenous probability we consider the optimal number of random audits, alternative audit strategies, and the choice between audit types (“hard”, where all concealed income is revealed at a higher cost of audit, or “soft”, where only part of concealed income is revealed, but at a lower cost). The focus of this section is on alternative audit strategies.

We analyze and compare the outcomes of four different audit strategies: random audits of the self-employed with a fixed probability (Fixed PA), audits of a fixed number of taxpayers in each self-employed occupation (Fixed NA), audits switching between self-employed occupations each period (Fixed NAA), and audits switching randomly between self-employed occupations (Fixed NAR). Rather than introducing the cost of audits explicitly, we construct the strategies with fixed numbers of audits to match the mean number of audits from the random audit strategy, so that on average over time the total cost of audits is the same for all four strategies. Given the same (average) cost, the best strategy is the one that delivers the largest amount of tax revenue (including the fines collected from caught evaders).

Figure 8 shows the amount of tax and fine revenues collected in every period for these four different audit strategies under the assumption of the target effect. The outcome is very similar for the bomb-crater effect. Although no strategy is uniformly better in every period, the strategy with the fixed number of audits for each occupation appears to deliver higher yield more often than the remaining two strategies. This observation




Figure 8: Tax and fine revenues under four audit strategies: target effect.

is verified in Figures 9 and 10, where the empirical cumulative density function (cdf) is plotted for all four series of tax and fine revenues from Figure 8 for the target effect and the corresponding figure for the bomb-crater effect (not shown here). Strategy Fixed NA dominates the other three strategies, in the sense of the first-order stochastic dominance.⁴ This implies that a revenue service with an objective function increasing in tax and fine revenue (in particular, the total amount of revenue) should prefer this strategy over the other three when maximizing the expected value of the objective function. This finding seems to be robust to the behavioural assumption on the taxpayers' immediate reaction to an audit.

7 Conclusions

The compliance decision combines a range of economic, psychological, and social elements. Included amongst these are perceptions of risk and attitudes toward risk-taking, the importance of social standing and conformity to group norms, and the transmission

⁴In this context, strategy A dominates strategy B in the sense of the first-order stochastic dominance, if for every level of revenue, R , the probability of collecting at least R is higher under A than under B . Equivalently, the empirical cdf of revenues collected under A is everywhere below (or to the right from) the empirical cdf of revenues collected under B .




Figure 9: Empirical cdf for tax and fine revenues under four audit strategies: target effect.




Figure 10: Empirical cdf for tax and fine revenues under four audit strategies: bomb-crater effect.

of information through social contacts. A compelling model of the compliance decision requires these components to be combined and embedded within a taxpayer equilibrium.

Agent-based modelling provides the ideal methodology for bringing disparate elements into a cohesive whole. The combination of the agent-based model with the structure of a social network to govern interaction provides a rich environment in which to explore compliance. A particular strength of agent-based modelling is that it has the potential to accommodate complex optimization and information updating processes.

The models that we have presented in this paper emphasize the importance of opportunities for non-compliance, and the link that this creates between occupational choice and risk attitude. Risky forms of self-employment will be chosen by those who are most willing to accept risk and to most fully exploit available opportunities for non-compliance. As a consequence, compliance behaviour can vary significantly across occupational groups.

The methodology is very flexible and is, therefore, able to incorporate recent advances in the theory of compliance. Our work emphasizes the role of attitudes, beliefs and opportunities, and draws ideas from advances in behavioural economics. A further advantage of an agent-based model of tax compliance is that it can incorporate a variety of different intervention strategies by the revenue service. We have contrasted random audits with three alternative strategies and have observed that the strategy with fixed number of audits in each occupation delivers the highest tax yield, keeping the average cost of audits constant across strategies. The strategies considered in this paper are not the only ones available to tax authorities. In particular, these strategies do not make use of the information obtained in the previous rounds of audit. One further direction of research is to explore the effect of *predictive analytics*, or the use of past information on taxpayers for predicting their future compliance behaviour, on audit outcomes.

Agent-based modelling is certain to become more influential in economic analysis as increased computing power permits ever greater model sophistication. Properly constructed

models will provide an ideal testing place for policy interventions that cannot be immediately tested in practice. Our models show a little of what can be achieved, but much more is possible.

References

- Allingham M., and Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of Public Economics*, 1, 323–338.
- al Nowaihi, A., and Dhami, S. (2007). Why do people pay taxes: Expected utility theory versus prospect theory. *Journal of Economic Behavior and Organization*, 64, 171–192.
- Andrei, A. L., Comer, K., and Koelher, M. (in press). An agent-based model of network effects on tax compliance and evasion. *Journal of Economic Psychology*. <http://dx.doi.org/10.1016/j.joep.2013.01.002>
- Bernasconi, M., and Zanardi, A. (2004). Tax evasion, tax rates and reference dependence. *FinanzArchiv*, 60, 422–445.
- Black, J., and de Meza, D. (1997). Everyone may benefit from subsidising entry to risky occupations. *Journal of Public Economics*, 66, 409–424.
- Bloomquist, K. M. (2004). Multi-agent based simulations of the deterrent effects of taxpayer audits. In Kalambokidis, L. (Ed.), *Proceedings of the 97th Annual Conference on Taxation* (pp. 159–173). Washington, DC: National Tax Association.
- Chateauneuf, A. (1994). Modeling attitudes towards uncertainty and risk through the use of Choquet integral. *Annals of Operations Research*, 52, 3–20.
- Davis, J. S., Hecht, G., and Perkins, J. D. (2003). Social behaviours, enforcement, and tax compliance dynamics. *The Accounting Review*, 78, 39–69.
- Guala, F., and Mittone, L. (2005). Experiments in economics: External validity and the robustness of phenomena. *Journal of Economic Methodology*, 12, 495–515.
- Hashimzade, N., Myles, G. D., Page, F., and Rablen, M. D. (in press). Social networks and occupational choice: The endogenous formation of attitudes and beliefs about tax compliance. *Journal of Economic Psychology*. <http://dx.doi.org/10.1016/j.joep.2012.09.002>
- Hashimzade, N., Myles, G. D., and Tran-Nam, B. (2013). Applications of behavioural economics to tax evasion. *Journal of Economic Surveys*, 27, 941–977.
- Kahneman, D., and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–293.
- Kanbur, S. M. (1981). Risk taking and taxation: An alternative perspective. *Journal of Public Economics*, 15, 163–184.
- Korobow, A., Johnson, C., and Axtell, R. (2007). An agent-based model of tax compliance with social networks. *National Tax Journal*, 60, 589–610.

- Myles, G. D., and Naylor, R. A. (1996). A model of tax evasion with group conformity and social customs. *European Journal of Political Economy*, 12, 49–66.
- Pestieau, P., and Possen, U. M. (1991). Tax evasion and occupational choice. *Journal of Public Economics*, 45, 107–125.
- Piolatto, A. and Rablen, M. D. (2013). Prospect theory and tax evasion: a reconsideration of the Yitzhaki Puzzle. *TARC Working paper 13-001*, Tax Administration Research Centre.
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior and Organization*, 3, 323–343.
- Slemrod, J. (2007). Cheating ourselves: The economics of tax evasion. *Journal of Economic Perspectives*, 21, 25–48.
- Snow, A. and Warren, R. S. Jr. (2005). Ambiguity about audit probability, tax compliance, and taxpayer welfare. *Economic Inquiry*, 43, 865–871.
- Tesfatsion, L. (2006). Agent-based computational economics: A constructive approach to economic theory. In Tesfatsion, L. and Judd, K.L. (Eds.) *Handbook of Computational Economics* Vol.II (pp. 831–880). Amsterdam: North-Holland.
- Yaniv, G. (1999). Tax compliance and advance tax payments: A prospect theory analysis. *National Tax Journal*, 52, 753–764.
- Yitzhaki, S. (1974). A note on income tax evasion: A theoretical analysis. *Journal of Public Economics*, 3, 201–202.