Tax Compliance in a Simulated Heterogeneous Multi-agent Society

Luis Antunes[1], João Balsa[1], Paulo Urbano[2], Luis Moniz[2], and Catarina Roseta-Palma[3]

[1] GUESS/LabMAg/Universidade de Lisboa, Portugal
{xarax, jbalsa}@di.fc.ul.pt
[2] LabMAg/Universidade de Lisboa, Portugal
{pub, hal}@di.fc.ul.pt
[3] GUESS/Dinâmia/ISCTE, Lisboa, Portugal
Catarina.Roseta@iscte.pt

Abstract. We consider an individualised approach to agent behaviour in an application to the classical economic problem of tax compliance. Most economic theories consider homogeneous representative agent utilitarian approaches to explain the decision of complying or not with tax payment. However, a heterogeneous and individualised account of decision can be considered to explain certain apparently irrational behaviours. Ideas such as trust and peer perception may have a key influence in individual decisions, and thus transform the global results for society. In this paper, we apply the agent view of rationality to economic decisions and define a territory to be explored by agent technology and social simulations. We conclude that the multi-agent view can provide powerful results which might lead to significant economic policy implications.

1 Introduction

In this paper we extend our views on agent rationality and multi-agent based exploratory simulation to tackle a classic economics problem that continues to defy specialists: tax compliance. The real challenge is not to discover why people evade taxes, but why anyone would pay taxes at all, given the low fines and even lower probabilities of being audited. This is a challenge for a realistic account of rationality, a hard problem for decision theorists. Given that there seem to be some social motivations behind the individual decisions, we propose a multi-agent based exploratory simulation as a way to formulate new hypotheses and conjectures, and perhaps provide some decisive insights into this issue.

The main claim of this paper is that the agent approach to complex systems, enriched by an heterogeneous view of rationality can provide a more realistic picture of the agents true motivations. The idea is not only to improve prediction capabilities (a difficult endeavour because of the increase in complexity brought about by considering autonomous agents, instead of agents with homogeneous

J.S. Sichman and L. Antunes (Eds.): MABS 2005, LNAI 3891, pp. 147–161, 2006.
© Springer-Verlag Berlin Heidelberg 2006

rationalities), but also to get sharper insights into a hard and complex problem. These insights, conjectures and hypotheses to be tested enhance our deep knowledge of the issue, and may help find new direct and indirect actions to be taken by the policy makers that can be effective both at the global and at the individual levels.

The paper is organised as follows: in the next section, we briefly describe the individual decision model we use for our agents. The relationship between the individual mind (with its own independent, adaptive rationality) and the dynamics of the whole society is very important in this application scenario. The agents' perception of the use of the collected taxes, as well as other social issues as reputation, social trust in both the government and other agents, all have influence in the individual decision. So, in section 3, we present our methodological approach to the exploration of this experimental setting. In section 4 we present some standard economic theories about tax compliance. Section 5 presents a series of increasingly complex models of a tax-aware agent. We end the paper with some preliminary results and some conclusions.

2 An Individualised Approach to Rationality

We have investigated for some time the theoretical underpinnings of rational decision. The standing idea that the choice problem "has already been solved" by utility theory took a severe blow when the economics Nobel Prize was awarded to Kahneman/Smith, in 2002.

The problems with utility theory are well known in the literature, and have been acknowledged in the AI community since the 1950s [20]. However, it is still important to stress that Simon's objection to utility was that it would be in practice impossible to compute. If agents could be instantaneous calculators, utility would provide a *definition* of rationality.

What this view fails to account for is the power of the agent's *will*. Agents will act as they will, depending their own motivations. While Simon coined one of the most important phrases about agent rationality, "people have reasons for what they do" [19], he would still ideally see people as expected utility maximisers. Freedom of choice is the keystone of agent autonomy, but choice is not (only) about calculations or search, rather, it is about volition and motivation. Of course, it remains to be explained why agents choose in the way they do, as their motivations cannot necessarily all be internally generated, some of them are socially influenced individual constructions.

The notion of rationality we will use can be described as *individual*, *situated*, and *multi-varied*. Agent decision can be based not on a single measure of marginal utility but on a multi-dimensional notion of value, which provides a referential for decisions to be taken. This multi-dimensional notion of value provides a referential for decisions to be taken. Then, the consequences of these decisions are assessed, and adjustments are made to the choice mental mechanisms. This choice model is sufficiently adaptive to account for several decision problems, as was extensively tested through simulations [6, 4].

3 Social Simulation and the Micro-macro Link

As we hinted at in the previous sections, the social simulation paradigm might methodologically benefit from taking on the agent stance. In the particular issue of tax compliance, we believe that the agent paradigm will bring along the necessary individualist view of each agent, and so provide social heterogeneity. In the classical sociological accounts (and in particular, in neo-classical econometric theories), practitioners seem to be concentrated in finding a general, all-purpose law that rules everyone's decisions (in this case of complying or evading, and how much to evade). This is obvious in classical examples from the literature, such as Wintrobe and Gërxhani [23] or Andreoni, Erard and Feinstein al. [2].

When social simulation is carried out from the multi-agent systems (MAS) view (and it has been since at least the first SimSoc workshop [15]), the agent view on social interactions becomes foundational. The most important consequence of this is the necessity to take into serious consideration the micro-macro link as described by Conte and Castelfranchi [14]. Roughly described, this means that there is a non-trivial (in fact, complex) interaction between the agents' individual mind and related behaviour and the society dynamics.

In the tax compliance scenario, a simple example can be seen in the trust notion of Wintrobe and Gërxhani [23]. One of the notions of trust the agents possess is trust that the government will provide enough public goods with the collected taxes. According to this model, an agent will evade less the more she trusts the government to fulfil its part of the deal, and vice-versa. The complex view on this mutual dependency goes a little further in detail. What is decisive for the agent's individual decision is not the amount of money the government will invest in the well-being of its governees. Rather, it is the *belief* of the agent about that issue. The realisation of this may lead the government to invest more money in publicity of the popular measures, thus in fact reducing the amount of money available for those measures, and so lowering the overall well-being.

In previous work, we have addressed the issue of methodologically linking the agents' mental structures for decision with the design of the experiment itself [3]. The argument is that by trying to control the agents' and society's behaviour through influencing their motivational system (so respecting the agents' "cognitive and executive autonomy" as Castelfranchi puts it [11]) we obtain a more accurate picture of the real processes that govern individual and global behaviour. The purpose of this move is to be in a better position to interpret results and propose alternative trajectories. By fostering freedom of choice and preserving agent autonomy, we hope that we can influence simulations as little as possible. The ideal situation would be to have completely different people programming the simulations and evaluating their results. But the research cycle for exploratory simulation is complex, and keeping the same people in all stages in fact helps locate the most useful insights (see fig. 1).

We view the work described in this paper as yet another instance where the MAS stance can make a difference by providing a new approach to an old problem. However, the methodologies for MAS, and especially the simulation vein of MAS, can greatly benefit from being confronted with a real issue, and one

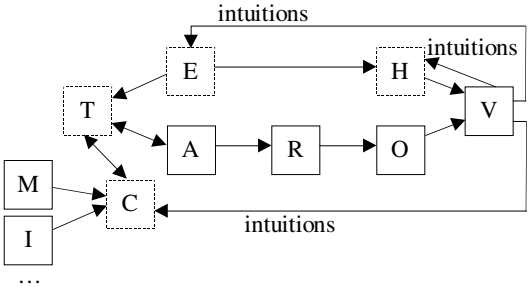

Fig. 1. Exploratory simulation. A theory (T) is being built from a set of conjectures (C), and in terms of the explanations (E) that it can generate, and hypotheses (H) it can produce. Conjectures (C) come out of the current state of the theory (T), and also out of metaphors (M) and intuitions (I) used by the designer. Results (V) of evaluating observations (O) of runs (R) of the program that represents assumptions (A) are used to generate new explanations (E), reformulate the conjectures (C) and hypotheses (H), thus allowing the reformulation of the theory (T). (From [3].)

that has been dealt with in other sciences, with different techniques, assumptions, and/or methods. In the long run, a better (more accurate, more robust, etc.) methodology for MAS experimentation is our ultimate goal, and these cases lead us one step forward.

4 The Economic Theory of Tax Compliance

Tax compliance has been studied in the economic literature from several different and complementary standpoints: public finance, law enforcement, organisational design, labour supply, and ethics (see the survey of Andreoni et al. [2]). Relevant notions in these approaches are equity; efficiency; incidence; effects of penalties; probability of evasion detection; organisational scheme design for taxation, audits and punishment; balance and structure of the labour market, among others.

The most compelling issue for us is the fundamental rational challenge posed to each individual, and how that relates to the overall behaviour of the society. Economists traditionally model individual tax evasion as if the individual is just adding one more risky asset to her household's portfolio [2]. Nevertheless, this theoretical approach fails to explain the behaviour that real societies display: households comply far more than could be predicted in this theory. For instance, in the USA, although fine value (or rate) can be neglected, and even though less than 2% of households were audited, the Internal Revenue Service (IRS) estimates that 91.7% of all income that should have been reported was in fact reported (numbers from 1988-1992-1995, cited from [2]).

In the next subsections, we will present several models of the choice problem each individual faces. Even though these models are very simple to start with, the predictive power they provide through analytical means is very limited, and so they already call for simulation so that the boundaries of their design can be properly explored. Since we intend to provide heterogeneous rationalities for our

different agents, and allow agents to change choice settings through interactions and other mechanisms, complexity becomes difficult to overcome. We examine these models as a search for the relevant ingredients to be used later on in more elaborate approaches.

4.1 Standard Theory

Allingham and Sandmo provide a seminal paper about tax evasion [1]. The terminology used there has become standard for the area. A taxpayer has exogenous income y, facing a tax rate t. The amount reported to the government is $x < y$, leaving $z = y - x$ unreported, and paying tax tx. The tax authority does not know the true income y, and has to enforce compliance through a policy of audits and penalties. The model goes on to assume that the enforcement policy is known to the taxpayer and depends on a probability p, with $0 < p < 1$. Further assumptions are that p does not depend on x, and that the tax authority is always able to discover the true value of y. Then, if θ is the penalty to be paid for every unit of income evaded, the cheating taxpayer will have to pay $\theta z + tz$. Given this, and assuming the taxpayer is risk averse, it can be shown that her expected utility if she decides to evade is [2, 23]:

$$(1-p)u[y(1-t)+tz] + pu[y(1-t)-\theta z)]$$

The effect of raising tax t for the "rational" taxpayer decision is not obvious. Raising t exerts contradictory forces in the agent's expected utility. Yitzhaki [24] observed that if the penalty is proportional to the amount of tax evaded (so θtz instead of θz), the model predicts that cheating will be reduced when the tax rate increases. Even if this settles our agents to rest about the decision, it hardly represents a breakthrough for the realism of the model, as noted in [23].

4.2 Public Choice Theory

The basic assumption of the 'public choice' field of economics is that citizens of democratic political jurisdictions perceive a connection between the taxes they pay and the government services they receive [23]. In face of our perception of the micro-macro link and individualised decision, this seems a promising alley to walk when looking for more complex models of taxpaying behaviour. The public choice view implies that every citizen knows that if taxes are reduced then a reduction in public services will follow. Even if the government is seen as a rational utility maximiser, the explanation for citizens to pay their taxes is that they trust the government to deliver the services it has promised. It is nevertheless rational for each individual to free ride, since the level of public services she receives from her own taxes will be negligible.

4.3 Trust Theory

The notion of trust, and related notions such as morality, ethics, categorical imperatives, has gained in recent years an increasing importance in the tax

evasion literature. Wintrobe and Gërxhani [23] mention two kinds of trust that influence the tax compliance decision: (*i*) trust that the government will provide the public goods it is supposed to provide and (*ii*) trust that fellow tax payers will indeed pay their taxes. Although the focus of Wintrobe and Gërxhani is on the first case, they argue that both types of trust would be identically correlated with the decision to comply or to evade taxes.

In Wintrobe and Gërxhani's trust model it is assumed that only one public good S exists, and all the citizens in this jurisdiction consume the same amount. Each citizen is also assumed to correctly calculate that the tax price to her of a unit of S is p. D is the demand, or marginal valuation curve for the public good. In the authors' analysis, they conclude that if taxes p rise without an accompanying increase in the amount of public good made available, then the propensity to evade taxes will increase. However, in certain conditions, if the rise in taxes is followed by a rise in the level of public goods provided, the propensity to evade will decrease. The role of trust becomes relevant in this model because the citizens will not have a direct perception of the real surplus in public goods deliverance. So, trust takes the role of this perceived surplus: citizens with a larger trust in the government (or that fellow citizens will pay, the argument holds just as well) will more easily support the government, which in turn will return their trust by trying to maximise the sum of citizens' surpluses in terms of public goods.

It is obvious that this trust relationship need not be symmetrical, as is implied by Wintrobe and Gërxhani. For instance, when addressing trust of type (*ii*), the authors mention (supported by Benjamini and Maital as well as Myles and Naylor [7, 17]) that theoretically there should be a "tipping point" when the number of tax evaders reaches a certain level and cause an epidemics of evasion. In fact, the overall behaviour of societies changes much slower, and all the real numbers show that people comply far more than any theoretical model predicts. It is also naïve to think that the government will always increase the public goods distributed, and only fails to do so because of corruption or lack of knowledge about people's desires. Central administration has costs, and these (as well as inefficiency) have always been neglected in all these models.

A more realistic notion of trust is used by Boadway, Marceau and Mongrain [10]. In this paper, the authors concentrate taxes associated with transactions. This determines that both participants in the transaction must agree whether or not to evade tax, and this joint decision requires that some trust bounds the two participants. Social interactions are modelled as a repeated iterated game, similar to the prisoner's dilemma. Agents are utility maximisers whose behaviour depends heavily on a personal parameter, θ, denoting *tolerance for dishonesty*, and trust is modelled as a kind of generalised *tit-for-tat* observed in all transactions. This being said, the conclusions of this research hold as far as its assumptions are met. Our contention is still that general laws are impossible to find that will govern everyone's mind. We *know* that some people will always pay their taxes, however unfair or nonrational that behaviour may be.

4.4 Extending the Classical Models

The survey paper of Andreoni et al. [2] is very complete and provides the cartography of the field. Although it divides models between "principal-agent" and "game-theoretical," there is a section when some moral and social issues are introduced. However, the effect of these issues on tax compliance literature are described as a "yet largely undeveloped area of research," and the authors add that "little is known or agreed upon about how best to include these effects in a theoretical or empirical analysis of tax compliance."

In light of this view, it is our contention that the agent-based approach with exploratory simulation is the answer to some of the problems raised by the unrealistic assumptions most analyses make. In particular, most models we have encountered carry the burden that all agents must follow the same utility patterns. With agent technology and individualised rationality, we can expect to find all sorts of agents following their own mind. With adaptive choice schemes we will provide agents with the ability to change the way they behave as a reaction to interactions and past events.

Among the proposals for extensions of the models in [2], we found the idea of trust, either in the government or in fellow taxpayers. The idea is to explain the empirically verified non-compliance with the sense of unfairness of the tax payed, either because of little service provided by the central authority, or as a comparison between what the agent pays and what she perceives others to pay. As we have seen, this idea has already been explored in subsequent papers, such as [10, 23]. However, the kind of analyses presented in these papers were also classic, and agent heterogeneity was not explored. Another extension suggested in [2] is the exploration of the moral obligation to be truthful, or the social consequences of being a know cheater. In mental terms, the senses of 'guilt' and 'shame,' respectively. This is another promising path for MAS techniques, since the mechanisms of reputation have been explored in detail for some time now (see the works of Castelfranchi, Conte and Paolucci [12] or Sabater and Sierra [18], for instance). Trust and deception is another area that has been dealt with within MAS for quite some time (see for instance, Castelfranchi and Tan's compilation [13]).

The problem of unrealistic assumptions of models may be a little more difficult to explore, because many of the parameters and constraints come directly from empirical observations or intuitions from practitioners. But even for that matter, the agent-based approach may be better suited and provide a quicker exploration of the space of possibilities, since through simulation the 'bad' choices can be detected and avoided sooner than if real empirical experiments were made.

4.5 Multi-agent Based Simulation

The first Multi-Agent Based Simulation model of income tax evasion was developed by Mittone and Patelli [16]. These authors followed the theoretical work of Myles and Naylor [17] that considers the existence of three categories of tax payers: honest, imitative and free riders. The behaviour of each tax payer category is characterised by a unique utility function. Utility is also influenced by

the public sector goods and services supported by tax contributions. This model was programmed using the SWARM package.

Bloomquist [8] developed a Tax Compliance Simulator in NetLogo in which taxpayers are represented as software agents that are defined by a set of characteristics like the income level, the income visible to tax authority, age, perception of enforcement activity, amongst many others. In [9], Bloomquist uses this simulator to conduct an experiment to estimate the deterrent effects of tax compliant alternatives.

5 A Model for Tax-Aware Agents

To commence the exploration of our simulation space we will adopt a very simple model that can fit the requirements of the game theoretical frameworks in the classical literature. However, it is important to stress that ours is not a game theoretical approach. In particular, we are not interested in studying the conditions for reaching some kind of equilibrium, rather we are curious about the overall dynamics of the system. This first approach to the tax-compliance social models will necessarily be very simple. We need to carefully study the dynamics of the simple model before we can start to introduce complexity (namely full agent autonomy and non-rational decision rules, such as moral imperatives).

Thus, in the rest of this section we will introduce an individual agent model that lays the bases for tax compliance decisions. Then we will introduce a model for the central authority, and its rules for tax collection, compliance enforcement, and finally redistribution. We made it a point to start where the classical econometric models start. However, now we depart from those models and delve into the richness of simulation. We propose new variables to illuminate the individual compliance decisions, then propose different categories of agents, each with its own set of motivations and decision procedures. We propose aggregate variables for the central authority and perception functions for the agents to observe them. We end up this section with a description of how agent interactions can determine streaks of behaviour that spread through the population and eventually change the overall dynamics.

5.1 The Ec0 Model

The first model we present is inspired in the classical models of [1], as described in [2]. An agent receives income y of which she will report $x < y$ and omit $z = y - x$. Reported income is subject to a tax t, so every agent will annually pay tx, if she chooses not to comply. In general, the tax may not be linear, so it could be $t(x)$.

Now we postulate a special agent, the central authority, in charge of enforcing compliance. Traditional rational agents would have no reason whatsoever to comply. So, the central authority uses audits and penalties to enforce tax compliance.

In the Ec0 model, we will assume that an agent is audited with probability p and if that happens the true income y will be found out. If the agent is caught evading she will have to pay a fine θ over the evaded amount, in total of θtz.

Note that some of the assumptions made at this point are very drastic:

- Audits are determined by a probability over every return filed. Note that even if probabilities may be a convenient *a posteriori* means for describing the returns examined, no tax agency will exclusively use rolling of dices to decide which return to audit.
- The probability of an agent being audited is independent of the past. This is unrealistic, since the tax agency will always audit an agent more if she has a history of evasion.
- By auditing the returns the tax authority will know the true value y of the income. Even nowadays, in the information era, tax agencies use all sort of data crossing among the millions of documents received.
- The probability of an agent being audited is independent of the probabilities of other agents being audited, and there is no limit for the number of audits to be carried out. As Andreoni at al mention [2], both these assumptions are wrong. More usually, agencies will have a budget which determines the maximum number of audits to be performed. On the other hand, as noted by Boadway et al. [10], agents perform (costly) efforts to avoid (or illude) audits, and this changes the probability that an evader is caught.
- The tax agency decision rule to audit a tax payer (based on probability p) is known by all the agents. This is also unrealistic, as most agencies indeed make enormous efforts to keep those rules secret.
- Another unrealistic assumption that changes the way agents should decide (for utilitarian agents, who compute expected utility) is the notion that when the tax authority finds an evader it will punish her only for the current year. More normally, the discovery of a cheater should induce further investigations on previous years (as well as prompt special attention for subsequent ones). So, the fine should enter the agent's calculations not as θz (or $\theta t z$) but rather as $\theta(z_\tau + z_{\tau-1} + z_{\tau-2} + \ldots)$. On top of that, the event of getting caught probably should influence the future probability of being audited (so increasing $p_{\tau+1}$). Obviously, this makes a substantial difference in both the individual choices and overall dynamics.

5.2 Ec1 and Ec2: Individuality and Adaptability

The first experiments with Ec0 revealed what could be analytically derived from the model: no one would pay taxes unless fines were unrealistically high. This result is typical of classical evasion models.

Our model Ec1 departed from Ec0 by postulating a new parameter to represent the "ethical stance" of the agent and allow for some evolution that could lead us to more realistic results. This ethical attitude is summarised in a number $\epsilon \in [0, 1]$. ϵ is used for the decision rule in such a way that when $\epsilon = 1$ the agent always complies to the tax due, and when $\epsilon = 0$ the agent always evades tax. Of course there are behaviours in between those two. For instance, in one experiment (see below) we used a rule that said the agent would pay when $\epsilon \geq \frac{1}{1+\theta}$, with $\theta > 0$ the fine.

A very important factor for the compliance decision in the classical models is the behaviour of agents in face of risk. Basically, an agent is risk neutral if her utility function is linear, and risk averse if the concavity of her utility function is facing downwards. Even if we are not using utility functions, agents will have to choose, and we can expect those choice functions to share some common ground with classic utility. The classical model assumes agents with decreasing risk aversion. This means that richer agents will more easily gamble with important quantities of money, and so would be more prone to evade. Of course here we have to distinguish between high income agents and wealthy agents. We introduce the quantity w of money an agent possesses. We expect w to be normally distributed, and update w every year by proposing a consumption rate γ to affect income y. To simplify at this early stage of research, we use $w_{\tau+1} = w_\tau + (1-\gamma)y'$, where $y' = (1-t)x + z$, the liquid income of the agent (we follow the notation introduced in section 5.1). Then the agent will be more prone to evade (so have greater ϵ) for greater values of w. For the experiments we used $\epsilon = \frac{w}{w+1}$.

Our Ec1 model credits agents with a limited amount of individuality. However, we need adaptivity, in order to have richer representation of true behaviours. Adding further degrees of complexity to the agents adds on enormously to the complexity of the whole society. So, we used the tactics of adding the mechanisms to allow one parameter to be dynamic, while freezing all others. A different possibility would be to experiment with that same parameter with several distributions, taking empirical studies as a starting point. The problem with that is that we do not know what kind of hidden assumptions were made in the models sustaining those studies.

For our Ec2 model, we proposed a dynamic behaviour for ϵ so had to find a rule for updating its value as a result of the consequences of past decisions [5]. Hence, model Ec2 includes the concept of memory of caught evasions and a factor (δ) for its degradation. If an agent is caught evading taxes, it will necessarily comply that year and will keep complying for some years onwards. So, if an agent is audited and found to be evading, we will set her ϵ to 1. Then, in subsequent years, we will update ϵ with the following rule: $\epsilon_{\tau+1} = \epsilon_\tau - \delta \in [0,1]$.

A more complex behaviour is obtained from relating this ethical attitude with the expectations of surplus received from the central authority in the public choice approach. If an agent observes a personal surplus of So and she was expecting Sp, variation update δ can be written as $1 - \frac{So}{Sp}$. This ensures that if $So > Sp$ then $\delta < 0$ and the agent moves towards paying taxes, and if $So < Sp$ then $\delta > 0$ and the agent moves towards evading.

The distribution of ϵ over the population should be based on psychological evidence, as well as the dynamics of its update. It should be more probable to find individuals with a tendency to always comply, and also individuals with a tendency to always evade. Even if some individual realises that he is paying too much and decides to start evading, it is not clear that this change would be gradual, as our model proposes. A possible credible distribution for ϵ could be similar to an inverted Normal, with lots of agents always complying (or slightly hesitating to do so) and lots of others always evading (or slightly hesitating to

do so). Our first experiments used a Uniform distribution, and we are currently undertaking experiments with other distributions.

5.3 Ec3*: Introducing Sociality

So far, the decision problem each agent has to face is an individual one. In Ec3* we introduce sociality: agents are aware of other agents and each individual decision depends on perceptions about the society as a whole. In subsequent models we will introduce interactions, and other forms of more complex social behaviours that contribute to overall social dynamics.

Following the idea of keeping experiments simple (in the path of constructing a "broad but shallow" agent design [21]), we introduce in the agent's design a perception of the level of overall fraud, $\Phi = \frac{\sum_{i \in Agents} z_i}{\sum_{i \in Agents} y_i}$. The agent's attitude towards this quantity will represent her level of comfort in face of unfairness. So we introduce $\phi \in [0, 1]$, the threshold of overall fraud admissibility for each agent. If $\Phi > \phi$ then the agent evades tax.

A slightly more complex model would include $\xi \in [0, 1]$ to represent the tolerance of agent in presence of global fraud. So, if $\xi = 0$ the agent ignores Φ and ϕ and complies, if $\xi = 1$ the agent follows the above rule.

Another interesting rule to explore is the local perception of fraud instead of the global perception. For instance, an agent will evade if a significant percentage of her neighbours evades. Quantity Φ is so changed to $\Phi_a = \frac{\sum_{i \in Neigh(a)} z_i}{\sum_{i \in Neigh(a)} y_i}$. Experimentation with these models is currently being conducted, and we expect that the crossing of its results with empirical data helps us make the necessary design options among the possibilities here outlined.

6 Preliminary Results

The results reported in this section were obtained conducting the described experiments using version 2.0.2 of the NetLogo framework [22]. NetLogo is a programmable modeling environment for simulating natural and social phenomena. It is particularly well suited for modeling complex systems developing over time.

In all experiments we consider two types of parameters: *internal* and *global*. Internal parameters are the ones that distinguish agents from one another (for instance, the agents' annual income). There are two types of global parameters: they can represent values that are equal for all agents (the tax rate, in the preliminary experiments) or values that the individual agents don't know at all (the probability of being inspected).

At this stage of research, we are still fine tuning some of the models and getting insights on how the several variables influence the system dynamics. So, in this section we will only hint at the kind of results we obtained, and present a brief analysis.

To test our model Ec2, we investigated how we could reduce the number of evading agents with the use of the memory effect on getting caught. So, for

Table 1. Percentage of evaders for several values of θ and p over δ

	θ	0.25			0.5		
	p	0.01	0.05	0.1	0.01	0.05	0.1
	0.01	.83	.51	.34	.75	.37	.23
	0.02	.91	.67	.51	.85	.55	.38
	0.03	.93	.74	.59	.89	.63	.46
δ	0.04	.95	.80	.67	.92	.69	.53
	0.05	.96	.83	.71	.93	.74	.59
	0.1	.98	.87	.77	.97	.84	.71
	0.2	.99	.91	.83	.99	.91	.83

reasonable (from the literature) values of the tax rate ($t = 0.30$), fine ($\theta = 0.25, 0.5$), and of number of audits to perform a year ($p = 0.01, 0.05, 0.1$), we varied the tendency for evading (the number δ to be subtracted from ethical attitude ϵ) from 0.01 to 0.05, and then 0.1 and 0.2. The results are shown in table 1. In every cell we have the mean percentage of evading agents in a population of 500 agents, with a normal distribution for the income and initial uniform ϵ subsequently updated as described in Ec2 above. We performed some dozens of experiments, and the numbers here presented are typical of what happened. Some stationary equilibrium was always reached, with some stochastic variability. Other measures were collected (total tax collected for redistribution, variability of ϵ, mean agent wealth, etc.) but their analyses are not completed yet.

As δ gets closer to zero, the memory effect of getting caught is stronger. So, agents will keep a good ethical attitude longer, and this gets reinforced by getting caught again. On the other hand, if agents do not evade taxes, their attitude leans towards evasion, at a pace set by δ. Looking at the table, we can see that with audit probabilities of 10%, we can still get relatively low numbers of evaders by keeping memory degradation slow. Without the memory effect, we could expect a mean of evaders approximately equal to 90%. These positive results are stressed by increasing the fine to 50% of the evaded amount. For instance, with a 1% probability of inspection we can have 25% of compliance, which is remarkable. We believe we can still improve these results by associating memory with the corresponding years, instead of considering the compliance decision of one year independent of all others.

7 Concluding Remarks

The main point of this paper is to argue that multi-agent systems concepts and techniques are a step forward in social science problems that have resisted more traditional approaches. The classical problem of individual tax compliance, as well as the problem of determining the correct tax enforcement policy, have constituted for decades a challenge for economists, sociologists and social psychologists. From our point of view, it is also an interesting problem for MAS practitioners, since it presents a clear case where the limits of situated

rationality are put to test, and the neo-classical economics approach of maximising expected utility remains wanting in face of the empirical results available. Tax compliance is also a great issue to test out agent based simulation methodologies and techniques, and to perform exploratory simulations that can help tackle the hot questions themselves, while gaining in experience and improving the necessary methodologies for experimentation with self-motivated agents.

In this paper we define the borders of this region we want to explore, and enumerate some of the difficulties already known in the literature, as well as criticize some of the assumptions made in other approaches. We then propose an approach based on multi-agent based simulation and describe agent models and experimental settings to be used. We already performed some experiments and the results are encouraging. However, this research is only beginning, and the complexity of the tasks at hand recommend prudence.

The promise of the first exploratory experiments, and perhaps the only true conclusion of this paper, is that the approach will deliver insights into the problem that were not available previously. Agent heterogeneity and individuality will provide a more realist account of the rational decisions that determine the overall behaviour of the society. And, above all, agent technology and exploratory simulation will (and already did) allow for the rehearsal of mechanisms to try out different design scenarios. The idea is to increasingly deepen the mechanisms and decision modules, in a plug-and-play manner, to provide enhanced adaptability to the world conditions, as well as allow for external (agent to agent, for instance by goal-adoption) and internal interactions (interacting, and possibly conflicting parts of the agent's mind).

Future work will first deal with the validation of empirical data, allowing us to settle on an accepted paradigmatic agent. Then, we will define this agent and continue exploration, by stretching and deepening some of its mechanisms. We will also look carefully into the measures to be taken from both experiments and real world, and the methods to validate the simulation results. Another important issue is to look into the role of the central authority in what concerns policy decisions. As well recognised in the literature, tax compliance (theory and policy) is a fundamental matter for today's countries, especially those in transition to democracy. We believe that more realistic multi-agent decision models will provide a helpful contribution in this policy area.

References

1. Michael G. Allingham and Agnar Sandmo. Income tax evasion: A theoretical analysis. *Journal of Public Economics*, 1(3/4):323–338, 1972.
2. James Andreoni, Brian Erard, and Jonathan Feinstein. Tax compliance. *Journal of Economic Literature*, 36(2), June 1998.
3. Luis Antunes and Helder Coelho. On how to conduct experiments with self-motivated agents (revised). In Mario Paolucci Gabriela Lindemann, Daniel Moldt, editor, *Regulated Agent-Based Social Systems: First International Workshop, RASTA 2002, Bologna, Italy, July 16, 2002, Revised Selected and Invited Papers*, volume 2934 of *Lecture Notes in Artificial Intelligence*, pages 31–47. Springer-Verlag, 2004.

4. Luis Antunes, João Faria, and Helder Coelho. Improving choice mechanisms within the BVG architecture. In Yves Lésperance and Cristiano Castelfranchi, editors, *Intelligent Agents VII, Proc. of ATAL'00*, volume 1986 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2000.
5. Luis Antunes, João Faria, and Helder Coelho. Choice: the key for autonomy. In Pavel Bradzil and Alípio Jorge, editors, *Proceedings of EPIA 2001*, Lecture Notes in Artificial Intelligence. Springer-Verlag, 2001.
6. Luis Antunes, Leonel Nóbrega, and Helder Coelho. BVG choice in Axelrods tributes model (revised). In Jaime Simão Sichman, François Bousquet, and Paul Davidsson, editors, *Multi-Agent Based Simulation II*, volume 2581 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 2003. Third Internacional Workshop, MABS 2002, Bologna, Italy, July 2002.
7. Y. Benjamini and S. Maital. Optimal tax evasion and optimal tax evasion policy: behavioral aspects. In W. Gaertner and A. Wenig, editors, *The Economics of the Shadow Economy*, pages 245–264. Springer-Verlag, Berlin, 1985.
8. Kim M. Bloomquist. Modeling taxpayers' response to compliance improvement alternatives. Extended abstract at the NAACSOS Conference, June 2004.
9. Kim M. Bloomquist. Multi-agent based simulation of the deterrent effects of taxpayer audits. 97th Annual Conference of the National Tax Association, November 2004.
10. Robin Boadway, Nicolas Marceau, and Steeve Mongrain. Tax evasion and trust. Technical Report 104, Center for Research on Economic Fluctuations and Employment, Université du Québec à Montréal, February 2000.
11. Cristiano Castelfranchi. Guarantees for autonomy in cognitive agent architecture. In Michael Wooldridge and Nick Jennings, editors, *Intelligent Agents: agent theories, architectures, and languages, Proc. of ATAL'94*, volume 890 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, 1995.
12. Cristiano Castelfranchi, Rosaria Conte, and Mario Paolucci. Normative reputation and the costs of compliance. *Journal of Artificial Societies and Social Simulation*, 1(3), 1998. http://www.soc.surrey.ac.uk/JASSS/1/3/3.html.
13. Cristiano Castelfranchi and Yao-Hua Tan, editors. *Trust and Deception in Virtual Societies*. Kluwer Academic Publisher, Dordrecht, 2001.
14. Rosaria Conte and Cristiano Castelfranchi. *Cognitive and Social Action*. UCL Press, London, 1995.
15. Nigel Gilbert and Jim Doran, editors. *Simulating Societies: the computer simulation of social phenomena*. UCL Press, London, 1994. Proceedings of SimSoc 1992.
16. L. Mittone and P. Patelli. Imitative behaviour in tax evasion. In B. Stefansson and F. Luna, editors, *Economic Modelling with Swarm*, pages 133–158. Kluwer, Amsterdam, 2000.
17. Gareth D. Myles and Robin A. Naylor. A model of tax evasion with group conformity and social customs. *European Journal of Political Economy*, 12(1):49–66, 1996. available at http://ideas.repec.org/a/eee/poleco/v12y1996i1p49-66.html.
18. Jordi Sabater and Carles Sierra. Reputation and social network analysis in multi-agent systems. In *The First International Joint Conference on Autonomous Agents & Multiagent Systems, AAMAS 2002, July 15-19, 2002, Bologna, Italy, Proceedings*, pages 475–482, 2002.
19. Herbert Simon. Rationality in psychology and economics. In R. M. Hogarth and M. W. Reder, editors, *Rational choice: the Contrast Between Economics and Psychology*. Univ. of Chicago Press, 1987.

20. Herbert A. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69:99–118, February 1955.
21. Aaron Sloman. Prospects for AI as the general science of intelligence. In *Prospects for Artificial Intelligence - Proceedings of AISB'93*. IOS Press, 1993.
22. U. Wilensky. NetLogo, 1999. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL. http://ccl.northwestern.edu/netlogo/.
23. Ronald Wintrobe and Klarita Gërxhani. Tax evasion and trust: a comparative analysis. In *Proceedings of the Annual Meeting of the European Public Choice Society – EPCS 2004*, 2004.
24. Shlomo Yizthaki. A note on income tax evasion: A theoretical analysis. *Journal of Public Economics*, 3(2):201–202, 1974.