Tactical Exploration of Tax Compliance Decisions in Multi-agent Based Simulation

Luis Antunes, João Balsa, Ana Respício, and Helder Coelho

GUESS/Universidade de Lisboa, Portugal
{xarax,jbalsa,respicio,hcoelho}@di.fc.ul.pt

Abstract. Tax compliance is a field that crosses over several research areas, from economics to machine learning, from sociology to artificial intelligence and multi-agent systems. The core of the problem is that the standing general theories cannot even explain why people comply as much as they do, much less make predictions or support prescriptions for the public entities. The compliance decision is a challenge posed to rational choice theory, and one that defies the current choice mechanisms in multi-agent systems. The key idea of this project is that by considering rationally-heterogeneous agents immersed in a highly social environment we can get hold of a better grasp of what is really involved in the individual decisions. Moreover, we aim at understanding how those decisions determine tendencies for the behaviour of the whole society, and how in turn those tendencies influence individual behaviour. This paper presents the results of some exploratory simulations carried out to uncover regularities, correlations and trends in the models that represent first and then expand the standard theories on the field. We conclude that forces like social imitation and local neighbourhood enforcement and reputation are far more important than individual perception of expected utility maximising, in what respects compliance decisions.

1 Introduction

Tax evasion is a serious problem in most economies, especially those in transition to democracy. Evasion undermines the central government budgets and expenditures, harming public welfare, and creates a sense of unfairness that can ultimately generate further evasion.

Interestingly, the scientific field that addresses tax evasion is known as *tax compliance* [2]. The decision to comply or evade is individual. When we consider rational individuals, who pursue their self-interest, we expect that the common behaviour would be to evade. However, in the real world, the numbers of compliance are quite high. Indeed, the literature of the field is mainly centered around discovering the adequate models to explain *why do people pay their taxes* [1,8,17,20]. Of course, central authorities would like to fully grasp the mechanisms underlying tax compliance and evasion, in such a way that they could ultimately promote evasion reduction [2].

Economists traditionally model individual tax evasion as if the individual is just adding one more risky asset to her household's portfolio [2]. Nevertheless,

L. Antunes and K. Takadama (Eds.): MABS 2006, LNAI 4442, pp. 80–95, 2007.

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 multi-agent systems, most accounts of agents assume limited rationality, that is, agents decide in such way to pursue their self-interest, based upon an idea of utility, and maintaining some degree of autonomy [9]. The deepest insight to approach the tax compliance issue by multi-agent based simulation comes from Simon's famous sentence "people have reasons for what they do" [18]. Each person/agent has her own limited rationality, and the notion of rationality here prescribed can be described as "multi-varied, situated, and individual" [4].

In this paper, we put together the methodologies behind multi-agent based simulation with rationally-heterogeneous agents and tackle the tax compliance problem. Our aim is to understand the mechanisms behind the compliance decision, both at the individual and collective (social interactions) level. Agent technology and exploratory simulation provide us with tools and methodologies that allow for the rehearsal of mechanisms to try out different design scenarios. Agent heterogeneity and individuality provide a more realist account of the rational decisions that determine the overall behaviour of the society.

Next section presents the broad context of this research. In section 3, we propose a hierarchy of models to explore expansions and alternatives to the standard theory. We describe how our models partially cover the design space, and propose a strategy for its exploration. Section 4 presents Ec_0, a model that represents the standard theories, which we use to introduce the concepts and terminology used on the field. In the following sections we introduce several enhancements to this basic model: Ec_0^τ, Ec_3^{*i}, and Ec_4^*. We examine their respective constraints, and report on experiments and simulations done over them. In section 9 we present the environmental setting we used, and in section 10 we conclude and discuss prospects for future work.

2 Context of Research

With the agent-view on computer intelligence, a lot of social issues gained relevance and built a huge source of metaphors and inspirations for societies of artificial agents. It was later that this collaboration between social scientists and computer scientists started to be fruitful for both sides. Since at least the first SimSoc workshop [12], most fields of social science began to endure the idea that computer agent societies could provide a powerful tool to conduct experiments in controlled environments in a principled way. Multi-agent based simulation was developed as a field where the inherently complex issues could be subject to controlled exploration, most of the times not to build or prove theories, but rather to find the "right" hypotheses, conjectures, intuitions, with which to carry on the scientific work [13,10]. The scientific questions to be answer are no longer only “what happened?” and “what may have happened?” but also “what are the necessary conditions for a given result to be obtained?” In exploratory simulation, the descriptive character is not a simple reproduction of the real social phenomena, and the prescriptive character cannot be simplistically resumed to an optimisation. To put things simply, the subject of research expands from the target phenomenon, and now includes its modelling, and not only real societies can be studied, possible societies can be studied as well.

In some recent papers ([6,3]) we argued that for characterising existing societies with enough realism as to allow solid explanatory power, and enough predictive power to permit policy recommendations, there exists the need for heterogeneous and adaptive rationality. In the tax compliance scenario, we have proposed some models of individual agents and of societies, and have been experimenting with them to gain insights into this complex issue, as well as set the grounds for theories that can be used both to explain individual and social behaviour, and to recommend central authority policies.

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 economics, public finance, law enforcement, organisational design, labour supply, and ethics [2], because it presents both theoretical and practical problems that are hard to be dealt with. 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. Because of its inherently complex social character, 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.

The idea of a society constituted by agents with heterogeneous rationalities is central in the research we are conducting. This view opposes the traditional endeavour of economics and particularly game theory, where all the agents follow the same general law, and societies are homogeneous in rationality and therefore, *qualitatively*, in trends for individual behaviour. This means that the sources of complexity in global behaviour are limited to circumstances of the world, and parametric features of the agent’s minds. Rather, with heterogeneous, self-motivated agents, societies are orders of magnitude more complex, since in every individual decision there is a potential for new, unpredictable behaviour. This is our bid for truly taking on the *open systems* challenge, as was proposed by Carl Hewitt in the 1970s [14,15].

Experimentation with such heterogeneous rational agents will foster different lines of research:

-i- Cognitive modelling: mind design is being experimentally challenged, especially in what concerns decision and motivation;
-ii- Multi-agent based simulation: social simulations where the central unit is the individual agent is a recent field, and with each new problem addressed,

more is discovered about the potential of this methodological approach to experimental social science. The issue of tax compliance is a rich research field for this approach. Exploratory simulation together with agent heterogeneity and individuality, will provide reasonability, realism, and the possibility to rehearse mechanisms and try out different scenarios;

-iii- Tax compliance modelling and policy: we can explore more accurate and realist models of the individual decision by the tax payer, and its consequences on the overall global society, particularly in what concerns policy decisions;

-iv- Methodology for experimentation with self-motivated agents: the validity of results obtained through simulation is always debatable, and it has been argued that self-motivation only makes the case worse [4]. By conducting simulations that span over a broad field of applications, analysing their results, and proposing theory refinements and agents' mind re-engineering, we will gather information that can inform a full-fledged methodology for experiments and simulations whose meaning can have an impact in the real target phenomena. This is especially important when there is the need to provide policy recommendations and expect their outcomes to be accurate.

3 The e*plore Methodology

The idea of using a collection of models to proceed with the exploration of the tax compliance problem has been used to illustrate a methodology for such problems. The base steps of this methodology come originally from Gilbert's lifecycle of simulation research [11]. The main ideas that go beyond those are centred around back and forth journeys to provide robustness and ensure exploration; progressive deepening of mechanisms in a broad but shallow design of agents, societies and experiments; and face complexity through exploration of model variability.

These are the steps of the e*plore methodology [5]:

i. *identify the subject* to be investigated, by stating specific items, features or marks;
ii. *unveil state-of-the-art* across the several scientific areas involved to provide context. The idea is to enlarge coverage before narrowing the focus, to focus prematurely on solutions may prevent the in-depth understanding of problems;
iii. *propose definition* of the target phenomenon. Pay attention to its operationality;
iv. *identify relevant aspects* in the target phenomenon, in particular, *list individual and collective measures* with which to characterise it;
v. if available, *collect observations* of the relevant features and measures;
vi. *develop the appropriate models* to simulate the phenomenon. Use the features you uncovered and program adequate mechanisms for individual agents, for interactions among agents, for probing and observing the simulation. Be careful to base behaviours in reasons that can be supported on

appropriate individual motivations. Develop visualisation and data recording tools. Document every design option thoroughly. *Run the simulations*, collect results, compute selected measures;

vii. return to step iii, and *calibrate everything*: your definition of the target, of adequate measures, of all the models, verify your designs, validate your models by using the selected measures. Watch individual trajectories of selected agents, as well as collective behaviours;

viii. *introduce variation* in your models: in initial conditions and parameters, in individual and collective mechanisms, in measures. Return to step v;

ix. After enough exploration of design space is performed, use your best models to *propose predictions*. Confirm it with past data, or collect data and validate predictions. Go back to the appropriate step to ensure rigour;

x. Make a generalisation effort and *propose theories and/or policies*. Apply to the target phenomenon. Watch global and individual behaviours. Recalibrate.

4 A Structure of Models to Explore the Tax Compliance Problem

The classical approach to tax compliance, as well as the main concepts and terminology of the field, are summarised in [2]. In [3], we have modelled the traditional agent in the income tax setting, and the corresponding society model: Ec_0. The details of this model are described in the next section.

A number of unrealistic assumptions were on the basis of Ec_0, and those could explain its little predictive power and accuracy in face of real world data. Significant changes were there proposed to the basic agent model of Ec_0, resulting in a series of models Ec_1, Ec_2, etc. Figure 1 depicts the structure of these models, and helps to illustrate the kind of trajectory we are attempting in our exploration of the space of possible designs. The idea behind this trajectory is to successively remove overcome the shortcomings of Ec_0, and use each new experiment to get a deeper insight into the problem and eventually converge on an appropriate model to face real data.

In the first models, the agent would resort to more complex reasoning to deliberate towards her compliance decision, by keeping track of past events, or

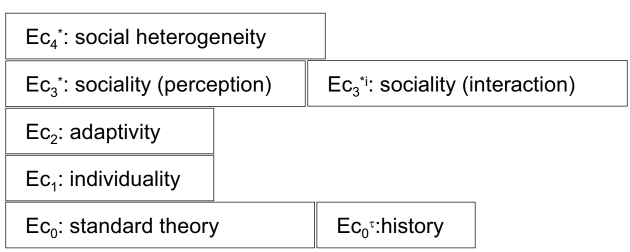

Fig. 1. Structure of designs for exploratory simulations in the tax compliance problem

accessing her individual characteristics. However, the decision was kept individual, as no social perceptions or interactions were taken into account in decision making.

In Ec_1 the agent possesses some wealth, and consumes it at a certain rate (γ). The agent's wealth determines its propensity to comply or not to its tax duties. Agents also have their own tendency towards evading, which we crudely represented with a real number $\epsilon \in [0, 1]$. The decision to comply takes these individual features into account. Ec_1 is the first model where the agent can be said to have some kind of (limited) individuality.

In Ec_2 we added some adaptivity. Ec_2 agents possess some evolution capabilities, namely, their tendency towards evading (or not) is dynamic, and evolves towards non-compliance at a given rate (δ). Some times agents are caught evading, and this tendency to non-compliance is completely canceled for a period of time we call memory, an individual agent parameter.

In Ec_3^* we introduced social perceptions as part of the data to be considered in decision. Agents have a global perception about the level of evasion in the society, and decide to comply or not according to their individual tolerance towards this perception.

The picture of the models to be considered is now completed with new models that encompass critical observations on the previous ones introduced in [3]. These new models and respective results will be introduced in subsequent sections of this paper, so here we only place them in the overall picture, as a way to deploy our strategy of exploration. In the remainder of this paper we will be concerned with tactical placement and setting of the models herein introduced.

So, in the next section we summarise the concepts, terminology and notations we used for Ec_0. Then we extend Ec_0 with history, introducing Ec_0^τ. This model slightly changes the standard theory to consider a criticism we produced in [3], that when an agent is caught evading the central authority will investigate not only the current year but also previous ones. We then present Ec_3^{*i}, a model where we explore the concept of imitation, and study how a core of stubborn agents can influence a whole society towards the behaviour they adopt, and how the distribution of two cores of stubborn agents with opposing behaviours can produce global effects. Finally, we introduce Ec_4^*, a model where two different breeds of agents are used to model tax payers and tax enforcers. Geography and locality become especially relevant in this latter model, as results are particularly dependent on initial conditions and performance circumstances of the simulation.

5 Ec_0: Modelling the Standard Theory

We now present model Ec_0, the basis of our pyramid of increasingly complex models. This model was introduced in [3] and we use it here to introduce the important concepts, terminology and notations, as well as to have a reference framework of ground results against which to compare our subsequent models.

The terminology used in [1] has become standard for the area. A taxpayer has exogenous income y, facing a tax rate t. The amount reported to the government is $x \leq y$, 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 additionally have to pay $\theta z + tz$. Given this, and assuming the taxpayer is risk averse, and that $u(\cdot)$ is the utility of money, it can be shown that her expected utility if she decides to evade is [2,20]:

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

For experimentation with Ec_0, we used a slightly simplified version of this formulae, concentrating only on what the agent saves by not complying. So, the evading decision is taken by each agent if the following inequality holds:

$$(1-p)u(tz) + pu(-\theta z) > 0 \tag{2}$$

The results produced by experimentation with Ec_0 are pretty much what could be expected from direct analysis of the decisions involved. Only in extreme and unrealistic conditions will the agents choose to comply. Observing that z (≥ 0) does not influence inequality 2, we conclude that tax payers evade when the following (all equivalent) inequations are satisfied:

$$p < \frac{t}{t+\theta} \qquad \theta < \frac{1-p}{p}t \qquad t > \frac{p}{1-p}\theta \tag{3}$$

Note that the decision is independent of the income value (y). Table 1 shows the turning point of individual decisions for usual values of the parameters. For instance, in the first section of the table we observe that for a tax of 30% and a fine of 50%, we need to inspect more than 38% of the tax payers to ensure compliance.

Table 1. Ec_0: Evasion point for usual values of θ, t, and p. Each cell contains, for each of the inequations 3, the point in which the truth value changes.

$\theta \setminus t$	0.10	0.20	0.30	0.40
0.25	0.29	0.44	0.55	0.62
0.50	0.17	0.29	0.38	0.44
$p \setminus t$	0.10	0.20	0.30	0.40
0.01	9.90	19.80	29.70	39.60
0.05	1.90	3.80	5.70	7.60
0.10	0.90	1.80	2.70	3.60
$\theta \setminus p$	0.01	0.05	0.10	
0.25	0.0025	0.0132	0.0278	
0.50	0.0051	0.0263	0.0556	

As expected, the numbers that promote compliance are quite stressful. For instance, for a tax rate of 40% and with a fine of 50%, the central authority would have to inspect over 44% of the tax returns to ensure overall compliance. Or, for an average tax rate of 30% and a probability of inspection of 1%, the fine would have to be above 2970% over the evaded amount. Or still, for a fine rate of 50% and a probability of inspection of 1%, the tax rate would have to be less than 0.51% to encourage compliance. None of these values is the least reasonable in face of what happens in real life. It remains to be captured by the model *what* leads people to comply.

6 Extending Ec_0 with History

One criticism that was quite prominent in [3] was that it would be rather awkward for the central authority to discover an evader agent and not investigate previous years. In fact, the common practice is that once a tax payer is inspected once, not only she will be investigated for her past, but also she will continue to be investigated in the future, even if she did not evade at all (although we are not considering the future in the model).

Model Ec_0^τ incorporate the previous history of the agent in her own utility calculations. Instead of considering θz, the amount of penalty will be $\theta(z_\tau + z_{\tau-1} + z_{\tau-2} + \ldots + z_{\tau-n})$, where τ is the current year. In most European countries, tax reports can be scrutinised reaching 10 years back, so we picked $n = 9$.

If, on top of this, interests are charged on past due taxes and/or fines, the compliance equation above significantly changes, and could indeed produce the tax compliance behaviour we observe on most Western countries. To simplify, we show only calculations for the simplified compliance rule:

$$(1-p)u(tz_\tau) + pu[-\theta \sum_{i=0}^{n} (z_{\tau-i}) - \sum_{i=0}^{n} ((1+r)^i t z_{\tau-i})] > 0 \qquad (4)$$

where r is the going interest rate for delayed payments. Note that the decision is taken in year τ, but only the 10-year aggravated penalties need be included in the decision, as the gains were already taken into account in the previous year's decisions. The decision to evade was already taken, there is nothing to be gained in year τ about that money, only to be risked. There are also some more simplifying assumptions, for instance, the interest rate is fixed over the years, as well as the fine rate, the tax rate and the probability of inspection. In a stable economic setting these options do not distort in any way the results of simulation, but in unstable settings, refinements must be made for the sake of realism.

In table 2 we can see the point in which the decision of whether to comply or not changes, for fixed usual values of the parameters. In the top part of the table, we fix some values for the tax and fine rates, and point out the smallest value for the probability of inspection that ensures full overall compliance. We note that there is a substantial decrease in the percentage of tax returns to be

Table 2. Ec_0^{τ}: Point in which the decision to comply changes for usual values of θ, t, and p

$\theta \setminus t$	0.10	0.20	0.30	0.40
0.25	0.23	0.31	0.36	0.39
0.50	0.15	0.25	0.28	0.31
$p \setminus t$	0.10	0.20	0.30	0.40
0.01	9.81	19.61	29.41	39.21
0.05	1.81	3.61	5.41	7.21
0.10	0.81	1.61	2.41	3.21

audited. For instance, for an average tax rate of 30% and a fine of 25%, we pass from 55% (in table 1) to 36% tax returns to be audited in order to ensure full compliance.

In the bottom part of table 2, we fix usual values for the tax rate and the probability of inspection, and we observe the smallest value of the fine rate that ensures full overall compliance. Here, the decrease in the fine rate is proportionally very small. It is obvious that any individual taxpayer is indifferent between a fine rate of 570% and another of 540% when deciding about compliance.

We conclude that these modifications alone do not have a very significant impact in the behaviour of tax payers, much less the overall behaviour of the society. For any reasonable values of p, t and θ, percentages of complying agents are quite small and far from reality.

Table 3. Evolving the percentage of evaders by changing θ, p, and n simultaneously

	$n = 10$		$n = 15$
	$p = 0.01$	$p = 0.02$	$p = 0.01$
$\theta = 0.5$	85%		
$\theta = 2$		24%	
$\theta = 4$	73%	5%	6%

However, when experimenting with increasing the probability of inspection p by very small amounts, we noticed that everything else being equal, the impact on the number of compliers of passing from $p = 0.01$ to $p = 0.02$ was far more significant than if we passed from, say, $p = 0.1$ to 0.2. So, we tried to manipulate more than one variable at the same time. In table 3 we have results that show some promise. By passing from $p = 0.01 \wedge \theta = 0.5$ to $p = 0.02 \wedge \theta = 2$, the number of evaders decreases from 85% to only 24%. More impressively, a very high fine rate of $\theta = 4$ will only conduct to 27% of compliers when $p = 0.01$, but for $p = 0.02$ it will yield 95% of compliers. And even for the smallest $p = 0.01$ (we cannot forget that it is expensive to conduct audits, far more than to increase fines), we can achieve 94% of compliance by examining tax returns back 15 years, instead of 10. The problem here would be that a simple strategy of 'killing' companies and founding new ones with the same assets could have the effect of 'cleaning up' the dirty past and getting on.

In any case, there remains an important gap to be covered by our models. We should note nonetheless that for any agent decision model M we can produce the corresponding history-aware model M^τ and enhance our coverage. However, the promise of Ec_0^τ is that perhaps there are still individual decision mechanisms to be explored in the classical theory and that the social front is not the only one to be considered. We will surely conduct further investigations along these ideas.

7 Stubbornness and Imitation

There is a considerable amount of literature about formation of consensus among homogeneous and heterogeneous agents. In multi-agent systems, the issue has been studied by Kaplan [16], and later by Urbano [19]. Urbano showed that a mechanism of imitation together with a small percentage of stubborn agents (agents that would not change opinion whatever happened) could be enough to promote a global societal change. In this section we propose to adopt this mechanism to model Ec_3^* [3]. The resulting model is called Ec_3^{*i}. In Ec_3^{*i} our population is divided into three subgroups: the stubborn compliers (SC), the stubborn evaders (SE), and the imitators (I).

Agents travel randomly in a square grid, and meet with other agents. For the decision, the agent will follow Ec_2 rules, which consider the ethical attitude $\epsilon \in [0,1]$, such that whenever $\epsilon = 1$ the agent always complies and where $\epsilon = 0$ the agent always evades. So, the agent will pay her taxes if $\epsilon \geq \frac{1}{1+\theta}\frac{\omega}{\overline{W}}$, where ω is the wealth of the agent, and $\overline{W}$ is the average of wealth of all agents. If the agent evades and is caught, her ϵ will be updated to 1. In the opposite case, her ethical attitude decays by a quantity δ (regulating the memory of having been caught evading). The idea of using factor $\frac{\omega}{\overline{W}}$ amounts to consider that agents whose wealth is above the average will more easily risk larger amounts of money than poorer agents. Agent's income is consumed at a rate γ, the remaining amount is added to her personal wealth.

Along the spirit of Urbano's investigation, we propose very simple and immediate mechanisms for imitation. The agent's individual attitude is publicly known (we will remove this constraint later on, and investigate on how the agent's behaviour can be affected by the reputation it renders), and agents imitate others following one of these rules:

- :i: The agent looks at the other agents in the same square, and adopts as her new ethical attitude (ϵ) the average of their ethical attitudes.
- :ii: The agent looks at the other agents in the same square, and adopts the ethical attitude of the more committed agent (the one that potentially has the highest influence on others), that is, of the agent that has her ϵ closest to one of the extremes of interval $[0, 1]$.
- :iii: The agent keeps track of the previous n encounters, and takes on the average of the involved agents' attitudes;
- :iv: The agent only considers the other agents' ϵ if they have a higher energy (for simplicity's sake, say wealth).

Using rule i, we conducted simulations with typical values: $\theta = 50\%$, $p = 1\%$, and $t = 30\%$. Our aim was to find out the combination of reasonable features that could ensure some stable equilibrium with an acceptable amount of overall compliance.

After some preliminary runs, it was clear that in such a highly social setting, the model was very responsive to variations in population density. So, in table 4 we present, for imitation rule i, the amount of evasion found in the equilibrium state (about 2000 iterations of the simulation), considering no stubborn evaders, and different proportions of stubborn compliers. In the line labelled "% ev. sc" we have the outcomes for a scarcely crowded society, where the average imitation rate per iteration is around 3.8%. Line "% ev. mc" displays the same numbers for a medium crowded society, where the average imitation rate is 9.7%. Line "%ev. hc" represents a highly crowded society, and the average imitation rate is 23.7%.

Table 4. Ec_3^{*i} Rule i: Variation of evasion over different proportions of Stubborn Compliers for different population densities

Rule i	$t = 0.3$ $\theta = 0.5$ $p = 0.01$ $\delta = 0.01$ $\gamma = 0.99$ $SE = 0\%$							
SC	0	10	20	30	40	50	60	70
% ev. sc	74	63	52	45	34	29	22	16
% ev. mc	78	65	48	37	28	25	16	11
% ev. hc	86	60	35	20	15	11	4	2

Figure 2 depicts the results of these series of experiments. It is clear that if all the society is composed of only imitators (no stubborns), the main trend is to evade, and agents imitate each other, hence reinforcing that trend. The introduction of a small proportion of stubborn compliers (say 20%) produces substantial effects in compliance for any density of population. In general, when we increase the relative number of stubborn compliers, the corresponding increase in compliance is always greater by a significant amount. This effect is particularly dramatic in highly crowded societies.

We conducted more experiments with the different imitation rules (ii, iii, and iv), but the results were not as exciting. An important parameter in the experiments is δ, the degradation of the memory of being caught evading. We remade the above experiments with $\delta = 0.10$ and the outcome was far worse. Agents seem to forget too soon that they were caught and start evading again. The consumption rate γ is quite high, about 99%. This means that agents do not save enough to be prepared to face a high penalty, and are better off complying. As happened in the previous series of experiments, increasing either the inspection (p) or the fine (θ) rates greatly reduces evasion. For instance, with $p = 10\%$ and $\theta = 100\%$, a stubborn compliers proportion of 20% yields an overall evasion rate below 5%.

From table 1 we can observe that for a fine of 50% and a tax rate of 30%, the minimum inspection rate to ensure total compliance is 38%. Here, with an

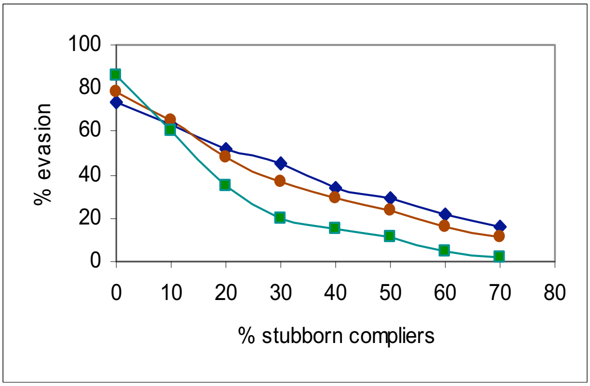

Fig. 2. Evolution of evasion when the fixed percentage of stubborn compliers varies. The square-dotted line represents a highly crowded society. The dotted line represents a medium crowded society. The diamond-dotted line represents a sparse society.

inspection rate of 1%, we can assure 65% of compliance with 20% of stubborn compliers, or even 89% of compliance with 50% of stubborn compliers.

This phenomenon we observed can explain why more people comply than standard theories predict. Stubbornness can be explained by personal characteristics, education, ethical stance, moral imperatives, social motivations, or even by some other mechanisms to be explored, such as reputation, imitation itself, altruism, fear, shame, political beliefs, etc.

The policy implications of these conclusions are substantial. If the stubborn compliance is appropriately encouraged (e.g. by offering prizes, such as tax reductions) and spread out, it will have a multiplying effect on overall compliance. In this case, it could pay off to bet on investing great efforts in building up a core of stubborn compliers, instead of dividing those efforts undiscriminantly over the whole population. The idea would be to *induce* or *favour* stubborn behaviour rather than *recognise* and *enhance* it, as it would be very difficult to distinguish a stubborn behaviour from an imitator with a very high threshold[1].

8 I Fought the Law and the Law Won

We now introduce an expansion of our previous models into Ec_4^*. In this model, we introduce a new breed of agents, the tax enforcers, or inspectors. An agent is only audited if she meets with one of these tax inspectors, and the inspector decides to audit her. This decision is taken autonomously, and so the whole concept of p ceases to be a number to be adopted blind by some anonymous central authority, and becomes an overall subjective goal of that authority, one that depends on individual decisions of the inspectors.

[1] We are thankful to an anonymous reviewer for pointing this out to us.

Geographical location (representing, more generally, complex social distances) becomes also an important issue, as well as trajectories that both tax payers and inspectors will go through. We keep trajectories random, using an uniform distribution to select among the eight candidates for each individual step. Later on, we will investigate mechanisms such as imitation of neighbours, clustering and flocking, to examine how these will influence the patterns of overall behaviour.

For tax-paying behaviour, we use the decision model of Ec_2. As to tax inspectors, their decision to inspect an agent is based on the following criteria. Central authority has a limited budget for auditing and inspectors, taken out of the whole amount of collected taxes. Each inspector receives a fixed amount of money per period, c_f, and each audit costs c_a. Inspectors are then assigned a personal budget b by the central authority. For now, these are obtained by dividing equally the overall budget B of the central authority by the number of inspectors.

When deliberating about whether or not to inspect an agent i, the inspector considers how much money he has got left from his budget, how much the audit will cost, and how much due tax and fine the audited taxpayer is expected to provide. These calculations are based on individual experience (frequency of successes in previous audits) and on the wealth of the inspected agent. If the information is available about previous evasions, the inspector can take that into account. For this purpose, inspectors exchange information among them about previous caught evaders. This exchange happens only when inspectors meet with other inspectors in their trajectories. Machine learning techniques can also be used to improve inspectors' performance.

With this inspection policy, we eliminate further criticisms of Ec_0 [3]: that audits are determined by a probability; that the probability of an agent being audited is independent of the past; that the probability of an agent being audited is independent of the probabilities of other agents being audited; that the cost of an audit is irrelevant and there is no limit for the number of audits to be carried out.

Experiments with Ec_4^* are reported in [7]. Our findings show that the overall compliance behaviour can be quite high in this new setting, given the appropriate fine tuning. Moreover, the distribution of the prerogative to audit tax payers from the central authority to autonomous inspectors can help meet the conditions for compliance in other models. However, individual decision mechanisms should still be enhanced, possibly through the use of context dependent adaptive functions. On top of that, the complexity and multiplicity of social distances must be taken into account in the simulation.

9 Experimental Environment

The experiments here reported have been programmed in NetLogo 3.0, of the Center for Connected Learning and Computer-based Modelling of the Northwestern University (Illinois). Figure 3 is a screen shot of our application running an experiment with Ec_3^{*i}.

The typical simulation would have a square world with 100×100 slots, where 500 agents would evolve. Then it would run until some equilibrium was found, which would happen at most around iteration 200 for the first experiments and around iteration 2000 for the simulations of section 6.

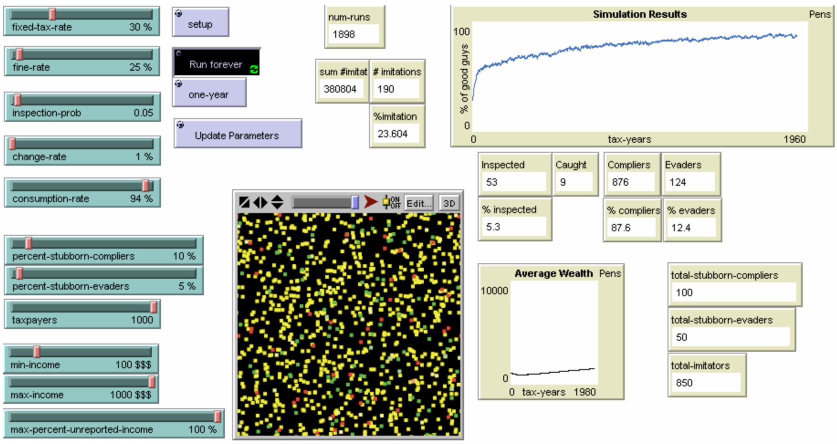

Fig. 3. A snapshot of application Ec^* running model Ec_3^{*i}

Development under the NetLogo 3.0 environment was gentle and swift, and the collection of data was straightforward. The results we present in the paper are mostly taken from typical runs, although sometimes we run several experiments and averaged the result. The exploration of the model parameters to obtain some desired outcome was conducted mainly by setting usual values for some parameters while thoroughly spanning over the remaining ones. In order to find which are the key parameters, as future work we intend to conduct a sensitivity analysis.

10 Conclusions

Tax compliance is a challenge posed to each individual in a society, and one that defies the standing proposals for explaining rational behaviour, as well as (re)produce it in artificial agent experimental settings. When the issue is looked upon from a social standpoint, experimental complexity grows steeply and comprehension of the observed phenomena becomes harder and even more defying.

We take on this challenge from the multi-agent based simulation standpoint. Our proposal is to consider situated, multi-dimensional, adaptive and individual decision models as a means to provide a richer representation of the actors in a simulation. In this paper, we report some exploratory simulations that have allowed us to have a deeper insight of the mechanisms involved in both the individual decision and the dynamics of societal behaviour.

We propose models that promote the explanation of compliance mechanisms to a level that the classical theories could not achieve. This is the first step towards a robust approach to tax compliance, able to predict real outcomes, and eventually to propose policy recommendations, for central authorities. This is particularly important in countries where public deficits and tax evasion are very high. Among these models, we introduced historical inspections and imitation behaviours, and obtained particularly good results. Stubbornness and leadership can nurture a stable equilibrium with good overall compliance levels, provided the appropriate constraints are assured.

Future work will focus on the completion of the series of experiments, as well as calibrating some new models and mechanisms. An important step towards this calibration will be to conduct a thorough sensitivity analysis on the models we already have run in our simulations: it is important to explore the parameter space and investigate whether we can be located in special niches or particular locations in that space. In particular, we will investigate how our models are influenced by the particular distribution of income we use to model our agents.

We also want to obtain real empirical data to help this fine tuning task, as well as validate our present results and design options. We also plan to compare our tax simulations with other compliance behaviours, such as public transit. Another idea is to stretch out the kind of choice functions we use, and consider prospect theory, infinite penalties, moral imperatives, and other individual sources of decision. Mechanisms for spreading of reputation and social stigma will also be investigated.

Acknowledgements

The authors wish to express their gratitude to the anonymous reviewers who throughout successive revisions of this paper contributed to its clarification and improvement.

References

1. Allingham, M.G., Sandmo, A.: Income tax evasion: A theoretical analysis. Journal of Public Economics 1(3/4), 323–338 (1972)
2. Andreoni, J., Erard, B., Feinstein, J.: Tax compliance. Journal of Economic Literature 36(2) (1998)
3. Antunes, L., Balsa, J., Moniz, L., Urbano, P., Palma, C.R.: Tax compliance in a simulated heterogeneous multi-agent society. In: Sichman, J.S., Antunes, L. (eds.) MABS 2005. LNCS (LNAI), vol. 3891, Springer, Heidelberg (2006)
4. Antunes, L., Coelho, H.: On how to conduct experiments with self-motivated agents. In: Lindemann, G., Moldt, D., Paolucci, M. (eds.) RASTA 2002. LNCS (LNAI), vol. 2934, Springer, Heidelberg (2004)
5. Antunes, L., Coelho, H., Balsa, J., Respício, A.: e*plore v.0: Principia for strategic exploration of social simulation experiments design space. In: Proceedings of The First World Congress on Social Simulation, Kyoto, Japan (2006)

6. Antunes, L., Del Gaudio, R., Conte, R.: Towards a gendered-based agent model. In: Proceedings of Agent 2004 Conference on Social Dynamics: Interaction, Reflexivity and Emergence, Chicago, USA (2004)
7. Balsa, J., Antunes, L., Respício, A., Coelho, H.: Autonomous inspectors in tax compliance simulation. In: Proceedings of the 18th European Meeting on Cybernetics and Systems Research (2006)
8. Boadway, R., Marceau, N., Mongrain, S.: Tax evasion and trust. Technical Report 104, Center for Research on Economic Fluctuations and Employment, Université du Québec à Montréal (February 2000)
9. Castelfranchi, C.: Guarantees for autonomy in cognitive agent architecture. In: Wooldridge, M.J., Jennings, N.R. (eds.) Intelligent Agents. LNCS, vol. 890, Springer, Heidelberg (1995)
10. Conte, R., Gilbert, N.: Introduction: computer simulation for social theory. In: Gilbert, N., Conte, R. (eds.) Artificial Societies: the computer simulation of social life, UCL Press, London, UK (1995)
11. Gilbert, N.: Models, processes and algorithms: Towards a simulation toolkit. In: Suleiman, R., Troitzsch, K.G., Gilbert, N. (eds.) Tools and Techniques for Social Science Simulation, Physica-Verlag, Heidelberg (2000)
12. Gilbert, N., Doran, J. (eds.): Simulating Societies: the computer simulation of social phenomena. Proceedings of SimSoc 1992. UCL Press, London (1992)
13. Gilbert, N., Doran, J. (eds.): Simulating Societies: the computer simulation of social phenomena. UCL Press, London (1994)
14. Hewitt, C.: Viewing control as patterns of passing messages. Artificial Intelligence 8 (1977)
15. Hewitt, C.: The challenge of open systems. Byte (April 1985)
16. Kaplan, F.: The emergence of a lexicon in a population of autonomous agents (in French). PhD thesis, Université de Paris 6 (2000)
17. Myles, G.D., Naylor, R.A.: A model of tax evasion with group conformity and social customs. European Journal of Political Economy 12(1), 49–66 (1996), http://ideas.repec.org/a/eee/poleco/v12y1996i1p49-66.html.
18. Simon, H.: Rationality in psychology and economics. In: Hogarth, R.M., Reder, M.W. (eds.) Rational choice: the Contrast Between Economics and Psychology, Univ. of Chicago Press (1987)
19. Urbano, P.: Decentralised Consensus Games (in Portuguese). PhD thesis, Faculdade de Ciências da Universidade de Lisboa (2004)
20. Wintrobe, R.: K. Gërxhani. Tax evasion and trust: a comparative analysis. In: Proceedings of the Annual Meeting of the European Public Choice Society – EPCS 2004 (2004)