Social Interaction in Tax Evasion

Lipatov, Vilen

Hannover University

28. March 2008

Social Interaction in Tax Evasion

Vilen Lipatov*

Hannover University, Königsworther Platz 1, 30167 Germany

Email address: *lipatov@mik.uni−hannover.de*

March 28, 2008

Abstract

We analyze the tax evasion problem with social interaction among the tax-payers. If the authority commits to a fixed auditing probability, a positive share of cheating is obtained in equilibrium. This stands in contrast to the existing literature, which yields full compliance of audited taxpayers who are rational and thus do not need to interact. When the authority adjusts the auditing probability every period, cycling in cheating-auditing occurs. Thus, the real life phenomenon of compliance fluctuations is explained within the model rather than by exogenous parameter shifts. Our analysis can also be applied to crime, safety regulations, employment and environmental protection, as well as other compliance problems.

JEL Classification: C79, D83, H26, K42

Keywords: tax evasion, learning, social interaction, behavioral rule, compliance

1   Introduction

The magnitude and importance of the shadow sector is hard to overestimate. Just to mention one case, the official estimate of the informal GDP for Russia is about 1/3

of formal GDP in the recent years[1]. A fundamental aspect of the informal activities is *tax evasion*, which is usually defined as an effort to lower one's tax liability in a way prohibited by law. The present paper focuses on this phenomenon, though tax avoidance and criminal activities can be analyzed in the same vein.

Most of the tax evasion literature to date is devoted to the income tax evasion. Such attention can be partially attributed to the existence of relatively reliable data on this matter (Tax Compliance Measurement Program (TCMP) in the US). Another reason might be tradition founded by the seminal model of Allingham and Sandmo (1972).

A detailed survey of the literature on income tax evasion can be found in Andreoni, Erard and Feinstein (1998). They identify two directions in the modeling of the strategic interaction between taxpayers and tax authorities: the principal-agent approach (e.g., Sanchez and Sobel (1993)) and the game theoretic approach (e.g., Reinganum and Wilde (1986), Erard and Feinstein (1994), Peter Bardsley (1997), Waly Wane (2000)). Both approaches treat the taxpayers as a single player maximizing her expected payoff from cheating. In particular, there is no direct interaction among the taxpayers in the literature we are aware of. In reality however a taxpayer is not an isolated decision maker; rather, she lives in a society and constantly interacts with other taxpayers.

This paper aims at characterizing the play in evasion game, when the taxpayers are boundedly rational and there is social interaction among them. Our taxpayers use a simple behavioral rule, and they decide whether to cheat or not cheat depending on the previous period behavior of themselves and of those whom they meet. This can be contrasted with more rational Bayesian updating, for instance when the agents have priors on the probability distribution of the auditing intensity and learn more about this distribution through their own play and interaction with others.

Our model is consistent with a number of stylized facts about evasion. First, in reality taxpayers possess poor knowledge of the audit rules, usually overestimating the probability of audit (Andreoni et al. 1998, pp. 844, 845). Accordingly, in our model they do not necessarily know it, but rather implicitly over- or underestimate it. When the number of interacting people is large, this setting also accounts for indirect channels of information transmission, such as media. Second, another feature

---

[1]A summary of attempts to estimate the size of tax evasion, avoidance and other informal activities is given in Schneider and Enste (2000). The results vary a lot with method and country considered; one common finding is that the shadow sector is growing over time.

of reality - the heterogeneity of information that taxpayers possess - is reflected in the initial distribution of choices between cheating and not cheating. Presumably, more informed taxpayers take the action that brings about higher expected payoff. Third, the-real world tax authority acts in a substantially different manner than an individual taxpayer. This organization has resources and incentives to gather a lot of the information, whereas every individual prefers not to incur the costs of information collection. The model reflects this asymmetry directly: the tax authority is updating its belief about the distribution of the taxpayers and accordinly adjusts it policy to maximize its revenue, whereas each of the taxpayers just follows a simple rule, being it imitation, learning, or payoff maximization. Fourth, tax evasion is an intertemporal decision, as claimed, for example, by the Engel and Hines' study (1999). In our framework the individuals have one-period memory that allows them to choose a strategy tomorrow on the basis of today's observation of the behavior of the others and their own. As the income reporting is a rare (annual) event, short memory can be a plausible assumption[2].

These four features are incorporated into a simple tax evasion model with an infinite taxpayer population and two income levels. We consider a large class of behavioral rules, including scenarios that vary according to the information taxpayers possess and their attitude towards punishment. At one extreme, taxpayers interacting in a small group know each other's income and evasion decision. At the other extreme, taxpayers interact in a large group, only observing caught evaders.

The main result of the model is the cycling dynamics of auditing and compliance. With non-committed tax authority, both the share of evading taxpayers and the auditing intensity of tax authority exhibit fluctuations giving rise to stable cycles. The system is cycling around an unstable steady state, in which the share of cheaters is the same as in the Nash equilibrium of the one-shot game, whereas the auditing probability is not related to its Nash equilibrium value. This happens because in the game the cheating is effectively determined by the rationality of tax authority, whereas the intensity of auditing is actually established by the learning rule and parameters of the game.

From the dynamics generated we can see that in presence of boundedly rational agents the equilibrium play does not actually occur. Therefore, the dynamics is

---

[2]A longer memory is in a sense the same as observing a larger number of agents and thus it can not alter the main qualitative findings.

necessary to be taken into account in order to make accurate inference about the welfare effects of various policies. The estimation of such effects, however, requires calibration of the parameters of the model, which is a separate issue.

Our model offers one potential explanation for a number of stylized facts. Firstly, non-zero cheating of audited taxpayers is obtained for the commitment case, which is certainly more plausible than the absolute honesty of the most of the principle-agent models (for example, Sanchez and Sobel 1993, Andreoni et al.1998). Secondly, in the non-commitment case, the following features of dynamics are explained: decreasing compliance (Graetz, Reinganum and Wilde 1986) and auditing probability (Dubin, Graetz and Wilde 1990, Adreoni at al. 1998, p.820) observed in the US in the second half of XX century. These patterns could not be explained by the literature to date since only static models have been used.

Additionally, an alternative explanation for the puzzle of too much compliance is offered. It is largely discussed in the literature that people comply much more than a simple lottery model of evasion predicts. Our results suggest that the reason might not be the presence of intrinsically honest taxpayers[3], but the fact that the system is far from the equilibrium. This is best illustrated in the commitment case: if the share of cheating taxpayers is converging to its equilibrium value from below, it looks as if taxpayers are cheating too little.

Our results are robust to a number of modifications. Firstly, more risk-averse taxpayers are equivalent to stricter learning rules. Secondly, with any finite number of income levels the dynamic patterns are preserved. Thirdly, the particular form of the learning rule does not matter for the most of qualitative features. Overall, it is only the bounded rationality of the taxpayers that is crucial for our results.

The rest of the paper is organized in the following way. Section 2 contains an outline of a simple static evasion game. A repeated version of this game with social interaction is analyzed in section 3, where we first state general results and then consider various learning rules and a numerical example. In section 4 we discuss applicability of our results to compliance problems other than tax evasion. Limitations of the model and possible extensions are mentioned in the concluding section.

---

[3]This is how the puzzle is usually resolved (for references see, e.g., Slemrod (2000)).

2 The model

2.1 A simple one-shot game

As a starting point for modeling the dynamics of evasion we take a simple one-shot game of tax evasion, based on Graetz, Reinganum and Wilde (1986). Intrinsically honest taxpayers (who can not evade for moral reasons) are eliminated from that model, as their presence does not change the results in the given setup. The timing is as follows:

1. The nature chooses income for each individual from two levels, high H with probability γ and low L with probability $1 - \gamma$;

2. Taxpayers report their income, choosing whether to evade or not;

3. Tax agency decides whether to audit or not.

It is obvious, that low income people never choose to evade, because they are audited for sure, if they report anything lower than L. At the same time, the high income people can evade, since with a report L the tax agency does not know, whether it faces a truthful report by a lower income taxpayer, or cheating from the higher income ones. The tax authority will never audit high income reports. Then the game simplifies to the one between higher income people and the tax agency:

	audit	not audit
cheat	$(1-t)H - st(H-L),\ tH + st(H-L) - c$	$H - tL,\ tL$
not cheat	$(1-t)H,\ tH - c$	$(1-t)H,\ tH$

where t is the income tax rate, s is the surcharge rate that determines fine for the amount of tax evaded, c is the audit cost; all these parameters are assumed to be constant and exogenously given for the tax-raising body[4].

It has been shown that under some mild parameter restrictions there is a unique mixed strategy Nash equilibrium $\left(q^{NE}, p^{NE}\right)$ with

$$q^{NE} = \frac{1-\gamma}{\gamma} \frac{c}{t(1+s)(H-L)-c}, \quad p^{NE} = \frac{1}{1+s}. \tag{1}$$

Simple comparative statics shows that auditing probability is decreasing in fine; evasion is increasing in costs of auditing and decreasing with fine, tax rate, income

---

[4]Endogenous determination of tax and penalty rate is an interesting task, but it constitutes the problem of a government rather than a tax authority. Moreover, it has been largely discussed in the literature, see, for example, Cowell (1990).

differential and share of high income people. Among other things, we have implicitly assumed here linear tax and penalty schemes, risk neutral individuals, and linear cost function for the tax authority. Even with such strong assumptions, repeated social interaction changes the predictions drastically.

2.2 Repeated social interaction

The importance of social interaction has for a long time been realized in economic profession. This, however, has not resulted in plethora of rigorous studies, as, on one hand, considering social interaction theoretically was too demanding in terms of computations due to intrinsic complexity of the matter; on the other hand, empirical studies on the topic are in most cases prohibitive due to simultaneity and reflection biases, as it is nicely outlined in Manski (2000). His paper actually provides a description of the state of the arts in the field of social interaction on the most general level.

According to Manski, interaction may occur in constraints, expectations, or preferences, as admittedly does everything in modern economics. Our paper considers constraint interaction: the constraint is the total taxable income raised by the taxpayers. More interestingly, the paper contains expectations interaction in the form of imitation by taxpayers. The preferences interaction occurs, if we allow the learning rules to be shaped endogenously, e.g. under the influence of social stigma.

Now consider the game presented in the previous section played every period from 0 to infinity. The populations of high income and low income taxpayers are infinite size with measures of γ and $1 - \gamma$ respectively, and this is a common knowledge. The true income of each taxpayer in a match is known to interacting taxpayers, but unknown to tax authority. The proportion q_τ of high income population is cheating by reporting low income at time τ. The agency is auditing the low income reports with probability p_τ, which is its private knowledge. Between the rounds the tax agency updates its belief about the distribution of taxpayers over cheating and not cheating, the high income agents are following a behavioral rule, potentially learning whose strategy performs better.

At time τ there are the following types of high income taxpayers: (i) honest, comprising proportion $1 - q_\tau$ of population and receiving payoff $(1 - t)H$; (ii) caught cheating, $q_\tau p_\tau$ of population with payoff $(1 - t)H - st(H - L)$; (iii) not caught cheating, $q_\tau(1 - p_\tau)$ of population with payoff $H - tL$. The tax agency is maximizing either

6

its long-run expected revenue by choosing auditing probability for all periods (commitment), or its expected revenue in the next period by choosing auditing probability for the next period (no commitment).

Note that we can not take a ready aggregate dynamics for the population of taxpayers because of the asymmetric nature of the players: in the no commitment case the tax authority is using myopic best response, whereas taxpayers imitate each other. Without such asymmetry, our game resembles emulation dynamics as it is defined by Fudenberg and Levine (1998), which is known to converge to replicator dynamics under some assumptions. However, these assumptions are not satisfied in our setup: most strikingly, each individual communicates with more than one other. Our aggregate dynamics thus does not converge to replicator dynamics and has to be derived for each imitation rule.

In our framework a behavioral rule for n people meeting[5] between the stages is a mapping from a set of outcomes yesterday $X = \{$honest audited, honest not audited, caught cheating, not caught cheating$\}^n$ into the set of actions today $Y = \{$honest, cheating$\}$. We only consider deterministic rules, so this mapping is a function. The evolution of the share of noncompliant taxpayers can be generally represented as

$$q_{\tau+1} = q_\tau + f(q_\tau, p_\tau). \tag{2}$$

The learning rule will shape the function $f(q,p) : [0,1]^2 \to [-1,1]$. In the appendix it is shown that this function is a polynomial of the order at most n in each of the arguments, and hence continuous. We shall call the behavioral rules that imply $f_q(q,p) < 0$ *stabilizing*, as the larger the share of noncompliance, the slower it is growing. Correspondingly, the rules with $f_q(q,p) > 0$ will be called *destabilizing*. The rules with $f_q(q,p) = 0$ will be called monotonic, as they imply monotonic dynamics for any feasible p, q.

Define $\hat{q}(p) : [0,1] \to [0,1]$ as a function that maps a set of possible auditing probabilities into a set of long-run outcomes of non-compliance shares, $\lim_{\tau \to \infty} q_\tau$, given $q_0 \in (0,1)$. Define also $\bar{q}(p) : \mathcal{P} \to (0,1)$ as a function that maps the set of auditing probabilities $\mathcal{P}$ into the set of interior steady state values of non-compliance shares. In other words, $\bar{q}(p)$ is an interior solution to the steady state condition: $f(\bar{q}(p), p) = 0$.

---

[5]Note that observing a *proportion* of the others in our setup does not make much sense, unless this proportion is infinitely small. The problem is that if a proportion of the whole population is met, a rule can condition the behavior on the own type of an agent, so paradoxically there is no social interaction.

For both $\hat{q}(p), \bar{q}(p)$ to be single-valued, we need to restrict our attention to the rules that produce a unique interior steady state. This will be true for almost all 2-person rules, all 3-person imitation rules, and many other rules with arbitrary number of interacting people.

The following assumptions seem reasonable and will be kept for the rest of the paper:

Assumption 1. $\bar{q}'(p) < 0$. Noncompliance is decreasing in detection, as otherwise the punishment is not perceived as such by imitation rule. We do not want to consider the rules that imply that our agents enjoy the fines.

Assumption 2. $\mu < 1 + s$, or $c < (1 + s) t(H - L)$. Otherwise the auditing is so wasteful that its costs are always higher than the benefits, so the tax authority will never audit.

The following proposition describes the evolution of the share of cheating taxpayers.

Proposition 1. Consider the dynamics of the share of evaders q for a behavioral rule that satisfies the assumption 1. For a stabilizing rule,

if $p_\tau \in [0, \bar{q}^{-1}(1)]$, then $q_{\tau+1} > q_\tau$;

if $p_\tau \in [\bar{q}^{-1}(0), 1]$, then $q_{\tau+1} < q_\tau$;

if $p_\tau \in (\bar{q}^{-1}(1), \bar{q}^{-1}(0))$, then $q_\tau < \bar{q}(p_\tau) \Longrightarrow q_{\tau+1} > q_\tau$, $q_\tau > \bar{q}(p_\tau) \Longrightarrow q_{\tau+1} < q_\tau$.

For a destabilizing rule the inequalities are reversed.

Proof. We shall prove the statement for the stabilizing rules, as it is completely analogous for the destabilizing ones. From (2) $q_{\tau+1} > q_\tau \Longleftrightarrow f(q_\tau, p_\tau) > 0$. For an interior solution ($p_\tau \in (\bar{q}^{-1}(1), \bar{q}^{-1}(0))$) we have $q_\tau < \bar{q}(p_\tau) \Longrightarrow f(q_\tau, p_\tau) > f(\bar{q}(p_\tau), p_\tau) = 0$, since $f_q(q, p) < 0$ (the rule is stabilizing). Correspondingly, $q_\tau > \bar{q}(p_\tau) \Longrightarrow f(q_\tau, p_\tau) < 0 \Longleftrightarrow q_{\tau+1} < q_\tau$. For a corner solution $\hat{q}(p) \in \{0, 1\}$. As $f_q(q, p) < 0$, either $f(0, p) < 0$ and $q_{\tau+1} < q_\tau$, or $f(1, p) > 0$ and $q_{\tau+1} > q_\tau$. Fix $q \in (0, 1)$. For $p | \bar{q}^{-1}(1) \leq p < \bar{q}(q), f(q, p) > 0$. By continuity $\exists \varepsilon > 0 | f(q, p - \varepsilon) > 0$. Now suppose $\exists p < \bar{q}^{-1}(1), q | f(q, p) < 0$. Then $\exists p_0 < \bar{q}^{-1}(1), q_0 | f(q_0, p_0) = 0$ that contradicts corner solution. Thus, $f(q, p) > 0$. The proof for $p_\tau \in [\bar{q}^{-1}(0), 1]$ follows the same lines.

We can see that for stabilizing rules with small (large) values of auditing probability the cheating is increasing (decreasing). Unexpectedly, in the interval of auditing probability values that result in interior solution, the change in the proportion of non-compliant taxpayers is negatively related to their number. This "anti-scale" effect is explained by the high enough detection probability, for which the caught cheaters contribute more to the increase of proportion of the honest, than the honest themselves.

The following remark will allow us to get a better feeling about the dynamic generated by the rules:

Remark 1. Iff a behavioral rule is stabilizing and $\exists \bar{q}(p), |f(q,p)| < |q - \bar{q}(p)| \forall q$, then for a fixed detection probability p the share of non-compliant taxpayers q is monotonically converging to $\bar{q}(p)$.

Remark 2. If a behavioral rule is stabilizing and $\exists \bar{q}(p), |f(q,p)| < |f(q + f(q,p),p)| \forall q$, then for a fixed detection probability p the share of non-compliant taxpayers q is converging to $\bar{q}(p)$.

The latter remark is a familiar contraction mapping; both remarks say that if the dynamics does not jump too vividly from period to period, it has to exhibit some convergence.

As for the tax authority, it maximizes its expected net revenue for any given learning rule in the population. Further we consider two cases for the behavior of tax authority. If it is unable to announce its auditing probability and keep it forever, we are in the "game theoretic" framework, and the our dynamics has two dimensions: already derived one for q and another one for p. We start, however, with a more simple case, when the auditors can credibly commit to a certain constant in time strategy (probability), and hence the dynamics is collapsing to one dimension.

2.2.1 Commitment

Assume that the authority commits to a certain auditing probability p once and forever (this corresponds to the principle-agent framework defined by Andreoni et al., 1998). This setup may seem unrealistic, but we have at least three reasons to consider it. Firstly, there is a well established tradition in the literature that deals with committed tax authority. Secondly, if the authority could choose whether to commit or not, it would commit, as this allows for a higher payoff. Thirdly, commitment

models are usually criticized due to the standard result of full honesty of audited taxpayers. As it will be clear from our results, such criticism does not apply in our framework.

Under commitment, the tax authority chooses p to maximize its steady state payoff

$$\gamma(1 - \hat{q}(p))tH + p(\hat{q}(p)\gamma(tH + st(H - L)) + (1 - \gamma)tL)$$
$$-cp(\hat{q}(p)\gamma + 1 - \gamma) + (1 - p)(\hat{q}(p)\gamma + 1 - \gamma)tL,$$

where a function $\hat{q}(p) : [0, 1] \to [0, 1]$ maps a set of possible auditing probabilities into a set of long-run outcomes of non-compliance shares, $\lim_{\tau \to \infty} q_\tau$.

The corresponding first order condition for interior solution is conveniently written as

$$\bar{q}(p^*) + p^* \bar{q}'(p^*) = \frac{q'(p^*) + \mu \frac{1-\gamma}{\gamma}}{1 + s - \mu}, \tag{3}$$

where

$$\mu := \frac{c}{t(H - L)}. \tag{4}$$

The condition (3) is a familiar equality of marginal expected benefit (left hand side) and marginal expected cost (right hand side) of auditing an additional taxpayer. The second order condition is then

$$2\bar{q}'(p^*) + p^* \bar{q}''(p^*) < \frac{\bar{q}''(p^*)}{1 + s - \mu},$$

where the notation is $\bar{q} = \bar{q}(p^*), \bar{q}' = \bar{q}'(p^*), \bar{q}'' = \bar{q}''(p^*)$. We keep this notation henceforth. We define $D := \bar{q}'' - (1 + s - \mu)(2\bar{q}' + p^* \bar{q}'')$. Note that $D \geq 0$ at the maximum, if $\mu < 1 + s$.

The comparative statics for the interior solution[6] gives the following relations:

$$\frac{dp^*}{ds} = \frac{\bar{q} + p^* \bar{q}'}{D}; \tag{5}$$

$$\frac{dp^*}{d\mu} = -\frac{\bar{q} + p^* \bar{q}' + \frac{1}{\gamma} - 1}{D}; \tag{6}$$

$$\frac{dp^*}{d\gamma} = \frac{\mu}{\gamma^2 D}. \tag{7}$$

We will call $\frac{\bar{q}(p^*)}{p^*}$ *average* steady cheating, and $\bar{q}'(p^*)$ *marginal* steady cheating.

---

[6]Note that these conditions do not hold on the border, e.g. for $p = \frac{2}{3}$.

Let us consider the forces at play behind the derivatives. The effect of the share of the high income taxpayers is most straightforward: it is always positive. Intuitively, more high income people make auditing more profitable. The effect of curvature embodied in the denominator is essentially the same in all three cases. From the second order condition it can be seen that for convex functions $\bar{q}''(p)$ we have $p^*(1 + s - \mu) < 1$, and for concave functions the opposite is true. Then both more convexity and more concavity increase denominator, thus reducing the effect of each parameter in absolute terms.

The higher fines bring about higher auditing effort, when the average steady cheating is larger than the marginal one (in absolute terms). In the opposite case, the tougher punishment relaxes the grip of the tax authority. Intuitively, the increased fine has a positive direct effect on both the effort of authority (makes detection more lucrative) and the share of compliant agents (makes compliance more attractive). The indirect effect is captured by the first term: increased compliance curbs the auditing. When non-compliance is relatively sensitive to auditing (and visa versa), the indirect effect is small compared to the direct one, so the increased fine results in lower detection probability.

For the normalized auditing costs, the direct effect is negative (the increase in costs makes auditing less profitable). The feedback effect is then positive, because more cheating calls for more auditing. Formally, the different direction of the effects compared to the case of fine is reflected in the different sign of the whole expression. There is an additional term $\frac{1}{\gamma} - 1$ that strengthens the direct effect.

To sum up, we arrive at the following proposition.

Proposition 2. With a committed tax authority and assumptions 1,2 satisfied, $\frac{dp^*}{d\gamma} > 0\left(\frac{dq^*}{d\gamma} < 0\right)$. If additionally $\bar{q} + p^*\bar{q}' < 1 - \frac{1}{\gamma}$, then $\frac{dp^*}{ds} < 0\left(\frac{dq^*}{ds} > 0\right), \frac{dp^*}{d\mu} > 0\left(\frac{dq^*}{d\mu} < 0\right)$. If $1 - \frac{1}{\gamma} < \bar{q} + p^*\bar{q}' < 0$, then $\frac{dp^*}{ds} < 0, \frac{dp^*}{d\mu} < 0$. If $\bar{q} + p^*\bar{q}' > 0$, then $\frac{dp^*}{ds} > 0, \frac{dp^*}{d\mu} < 0$.

Proof The proposition follows from explicit consideration of inequalities implied by equations (5)-(7) under assumptions 1,2.

The proposition actually states that the detection probability is increasing (and the steady state cheating non-increasing) under a very mild condition on convexity of the steady state relation between auditing and compliance. Given this condition, the auditing is affected positively by fines and negatively by costs, if the average steady

cheating is larger than the marginal one. The opposite is true if marginal steady cheating is substantially (how much is substantial depends on the share of high income taxpayers) larger than the average. The case in between gives probability negatively affected by both fine and generalized costs.

We can also see from the proposition that for some parameters the steady state noncompliance is decreasing in tax rate. This stands in contrast to the conventional result of the literature: the models featuring fine proportional to the tax evaded starting from Allingham and Sandmo (1972) result in higher tax rate contributing to compliance. In our model this does not hold as long as the indirect effect of the generalized cost outweighs the direct effect, which is in turn more likely with higher share of high income taxpayers and more steep steady state relation between auditing and compliance. This, as all other comparative statics results, will depend on the particular form of a learning rule.

The solution obtained can be compared with the Stackelberg-like equilibrium of the classical evasion game, when the tax authority moves first (much weaker asymmetry). Recall, that in this setup $q = 1$ if $p < \frac{1}{1+s}$, $q = 0$ if $p > \frac{1}{1+s}$, and undetermined for the equality. Since auditing is costly, the authority will choose either $p = 0$, $q = 1$, or $p = \frac{1}{1+s}$, $q = 0$. The latter is preferred whenever the auditing is not too costly, namely $c < \frac{\gamma}{1-\gamma}(1+s)t(H-L)$, or, in other terms, $\mu < \frac{\gamma}{1-\gamma}(1+s)$ (analogous to the expression in the dynamic version). Comparative statics is trivial in this setup: zero cheating result is independent of parameter changes as long as they do not violate rather mild condition of relatively not too expensive auditing. Auditing probability is decreasing in the surcharge rate, just as in the previous model. The solution of the static model is discrete, and the probability of audit jumps to zero for high enough μ or low s.

The prediction of the dynamic model appears to be more plausible, since non-zero cheating is not observed in reality. As it is known from the literature, the result of zero cheating in commitment case generalizes for more complicated models with continuum of taxpayers and presence of intrinsically honest taxpayers. Moreover, the commitment models are usually criticized on the basis of this unrealistic prediction. The model presented eliminates this fault, and allows us to reconsider the view of commitment as something implausible. Then it just boils down to the classical case of dynamic inconsistency, and the willingness to commit is equivalent to the planning horizon of the authorities.

An average taxpayer with high income in Stackelberg setting can only cheat or not cheat with probability one; in the dynamic case there is a possibility of a mixed equilibrium. This brings about higher "dynamic" payoff for the individual, if $p < \frac{1}{1+s}$, and lower payoff otherwise[7]. This simple result is straightforward: in the classical setup the equilibrium payoff of taxpayers does not depend on the auditing probability or the magnitude of fine. Hence, the expected payoff in the dynamic model is greater, if the audit probability is lower than in the static model, and visa versa.

2.2.2 No commitment

Under no commitment, the tax authority decides on the optimal auditing rule in every period, assuming that the distribution of the taxpayers has not changed from the last period $q_{\tau+1} = q_\tau$ (myopic best response). Then the expected revenue is determined by (??), where $\bar{q}(p)$ is substituted by q_τ.

The best response strategy is

$$BR(q_\tau) = \begin{cases} 0, & \text{if } q_\tau \leq \mu_1; \\ 1, & \text{if } q_\tau \geq \mu_1. \end{cases}$$

Here μ_1 is the level of cheating that induces switch of best response from zero to one or back:

$$\mu_1 := \frac{1-\gamma}{\gamma} \frac{c}{t(1+s)(H-L)-c}. \tag{8}$$

As the tax authority is very unlikely to jump from not auditing anybody to auditing everybody and back, we explicitly augment the choice of tax agency with inertia variable[8]:

$$p_{\tau+1} = \alpha BR(q_\tau) + (1-\alpha)p_\tau, \tag{9}$$

where α determines speed of adjustment; BR is the best response function, which is defined above as revenue maximizing p given the belief about the distribution of taxpayers. With $\alpha \to 1$, we are back to the case of jumping from 0 to 1 probability; with $\alpha \to 0$, the probability of audit stays very close to an initial level forever.

An interior steady state is described by the following pair:

$$q^{ss} = \mu_1, p^{ss} = \bar{q}^{-1}(\mu_1). \tag{10}$$

---

[7]Evaluating $I(\bar{q}(p), p) - I(0, \frac{1}{1+s})$, we get expression the sign of which depends only on the sign of $1 - p - ps$.

[8]the inertia assumption is common in the literature on learning and evolution - see e.g. Fudenberg and Levine (1998, p. 31). Convex cost of auditing would have a similar smoothing effect.

The following proposition shows that in the continuous time the system converges for a wide class of rules.

Proposition 3. Consider a learning rule characterized by transition function $q_{\tau+1} = q_\tau + f(q_\tau, p_\tau)$ with $f_i(q_\tau, p_\tau) < \infty, i = q, p$. With non-committed tax authority, the interior steady state is stable in continuous time, if the learning rule is stabilizing, $f_q(q, p) < 0$.

The proof is left to the appendix. Simulation results show that in discrete time setup the cycles around the steady state are observed for various rules. Intuitively, convergence a là remark 2 can not be achieved, if the detection probability is changing discretely. However close it comes to the steady state at a range of p that induces a certain direction of change in q, at some point p goes out of this range and q starts movement in the opposite direction - as a result we observe cycles.

Comparative static result for the steady state follow trivially from (8), (10) and Assumptions 1, 2:

$$\frac{dp^{ss}}{ds} > 0, \frac{dp^{ss}}{dt} > 0, \frac{dp^{ss}}{dc} < 0, \frac{dp^{ss}}{d\gamma} > 0.$$

Thus, the auditing probability in steady state is decreasing in costs of auditing and increasing in the share of high income taxpayers, the tax rate, the magnitude of fine, and the income differential. Compared to the Nash equilibrium, where probability to audit only depends on the surcharge rate, our result looks more plausible.

Still, for all parameters but s and γ the effects are the opposite of those in the commitment model. Whereas it is an open question what horizon a particular tax authority has, we can compare predictions of the two models by their conformability with stylized facts. First, it is a prevailing view that evasion is increasing in the tax rate (See, for example, Clotfelter (1983), Poterba (1987), Giles and Caragata (1999)), so here the commitment model seems to make a better job. Second, there is also a weak evidence that evasion is rising with the income (Witte and Woodbury 1985), and in this sense the long horizon authority is also superior. There is no convincing evidence on the influence of auditing costs on the auditing probability, and it is really difficult to say which model is closer to reality on this point.

In the rest point of the repeated game q is the same as in the Nash equilibrium of one shot game, $q^{ss} = q^{NE}$, since it is derived from the same maximizing revenue decision of tax authority. Auditing probability p^{ss} can be greater or smaller depending

14

on the parameter values and the learning rule. As for the payoffs, since q^{ss} is such that makes the tax authority indifferent between auditing and not auditing, its revenue is exactly the same in static and dynamic setups. With fixed q the payoff of an average high income taxpayer is unambiguously decreasing with p, so that comparison across the models is again ambiguous.
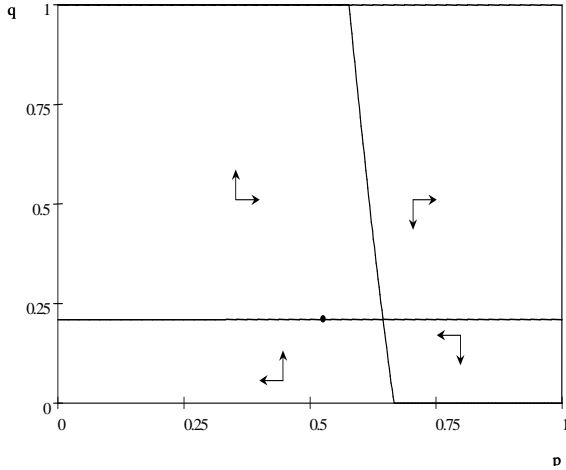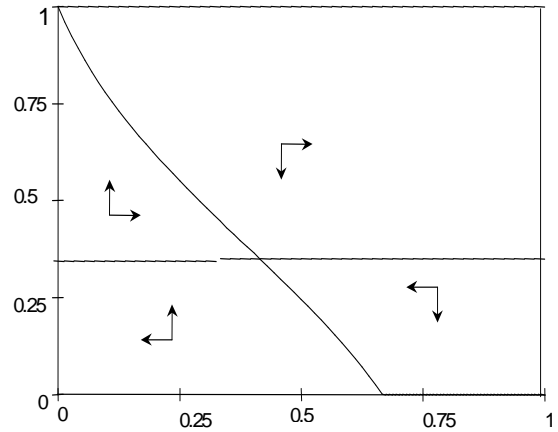

Figure 1. Phase diagram, average payoff rule

Figure 2. Phase diagram, effective punishment rule

On the phase diagram we can see that an interior steady state is the intersection of the horizontal constant auditing line $q = \mu_1$ and downward sloping constant compliance line $q = \bar{q}(p)$. Figures 1, 2 represent the dynamics generated by the rules considered in the next section. It is worth noting that the south-west and north-west parts of the picture is consistent with stylized facts presented in the introduction: both audit probability and the proportion of honest taxpayers decrease (second part of XXth century). According to this explanation, the observed behavior is out-of-equilibrium adjustment, and sooner or later the tax evasion will have to go down.

The southern part of generated dynamics also produces values of non-compliance lower than the Nash equilibrium. This can be taken as an alternative explanation to the puzzle of too high compliance, usually resolved by introduction of intrinsically honest taxpayers (Andreoni et al. 1998, Slemrod 2002).

To summarize our findings about the dynamics, we present the following corollary:

Corollary 1. Consider a stabilizing behavioral rule that does not "jump" too much in the sense of remark 2. The taxpayer population dynamics is characterized by

15

convergence to a steady state in case of committed tax authority and by stable cycles around a steady state in case of no commitment.

3   Particular behavioral rules

In our simple setting, any agent that has an option to not comply, may receive three distinct payoffs: from complying, from not complying and being audited, and from not complying and not being audited. We call the play of such an agent in the previous period its *type*, and thus we have 3 types, H (honest, or compliant), C (caught, or non-compliant and audited), N (not caught, or non-compliant and not audited). Here we describe the social interaction by considering behavioral rules that i) assigns the same action to all the agents in a match ii) assigns the same action to the members of matches characterized by the same composition of types. Further we consider more closely two rules characterized in the following table:

types met	share	average payoff	effective punish
HHH	$(1-q)^3$	honest	honest
NNN	$q^3 (1-p)^3$	cheat	cheat
CCC	$q^3 p^3$	honest	honest
HHN	$3 (1-q)^2 q (1-p)$	cheat	cheat
HHC	$3 (1-q)^2 qp$	honest	honest
HNN	$3 (1-q) q^2 (1-p)^2$	cheat	cheat
HCC	$3 (1-q) q^2 p^2$	honest	honest
NNC	$3 q^3 (1-p)^2 p$	cheat	honest
NCC	$3 q^3 (1-p) p^2$	cheat	honest
HNC	$6 (1-q) q^2 (1-p) p$	cheat	honest

Here the first row lists possible combinations of types met, the second reports corresponding shares of population, and the last two reflect the resulting play in the next period.

Having analyzed these two rules, we proceed by constructing their combination that allows for a change in attitude towards compliance depending on its popularity. Then we also examine a many people interaction rule, and conclude the section by a parameterization example.

3.1 Average payoff principle

Here we consider an imitation rule with very high informational requirements, that is not only the taxpayers who interact observe each other's income, but also each other's tax reports. Admittedly, this can only make sense if the size of interaction groups is small. In this case we can think of such groups as very close friends, which do share this information, for example, in CIS reality. Unfortunately, there are no formal studies on how widespread this phenomenon is, but based on personal experience, and experience of friends, and friends of the friends, the author believes that it is very relevant. Moreover, such information sharing may be relevant throughout the world, if we think of our interaction groups as families or siblings.

Collecting terms corresponding to "honest" from the table, we get

$$1 - q_{\tau+1} = (1 - q_\tau)\left((1 - q_\tau + q_\tau p_\tau)^2 + p_\tau q_\tau(1 - q_\tau + 2p_\tau q_\tau)\right) \tag{11}$$

This equation defines the aggregate dynamics of the population we were interested in. We see that this rule is stabilizing. Interior steady state level of cheating $\bar{q}(p)$ at given auditing is

$$\bar{q}(p) = \frac{2 - 3p}{1 - 3p + 3p^2}. \tag{12}$$

In the **commitment** case then the long-run non-compliance function is

$$\hat{q}(p) = \begin{cases} 1, & p \in \left[0, \frac{1}{\sqrt{3}}\right]; \\ \bar{q}, & p \in \left(\frac{1}{\sqrt{3}}, \frac{2}{3}\right); \\ 0, & p \in \left[\frac{2}{3}, 1\right]. \end{cases}$$

Let us call the probability that maximizes the authority's payoff (**??**) the optimal auditing probability p^*. Note that $p^* \notin \left(0, \frac{1}{\sqrt{3}}\right)$ and $p^* \notin \left(\frac{2}{3}, 1\right)$ because for constant q the objective function is linear in p. Furthermore, $p = 1$ is never optimal because the objective function is non-increasing on the interval $\left(\frac{2}{3}, 1\right)$. Hence, the only two possibilities for optimal p are $p^* = 0$ and p^* that satisfies the first order condition (3). The condition for corner solution ($p^* = 0$) is considered in the appendix, and it can be seen that corner solution results for values of γ small relative to values of μ. This is easy to interpret: with small share of high income people (γ) and low benefit from auditing (stH) it is better not to audit anybody, given that it is costly (c).

As for the speed of convergence, the average payoff rule is in a sense favorable to the cheaters: it takes a long time to approach no cheating equilibrium, and relatively

short time - all cheating one. Starting from the middle ($q = \frac{1}{2}$), getting as close as 0.001 to the steady state takes 597 periods for honesty case and only 14 periods for cheating case. This result was obtained by iterating the function $q_{\tau+1}(q_\tau, p)$ respective number of times. For the honesty case then $p = \frac{2}{3}$, $q_{597} = 0.001$; for the case of cheating $p = \frac{1}{\sqrt{3}}$, $q_{14} = 0.999$, whereas $q_1 = 0.5$ in both cases.

In the **no commitment** case the dynamics can be seen on the Figure 1. The variation in steady state p is very small: $\left(\frac{1}{\sqrt{3}}, \frac{2}{3}\right)$, compared to $(\frac{1}{2}, 1)$ in static case for $s < 1$. Hence, the difference in p for these two models is primarily dependent on s: for large values of fine Nash equilibrium gives less intensive auditing, and for small fines our model results in lower auditing.

3.2 Effective punishment principle

The idea of the effective punishment rule is that observing punished people (or being punished) is a sufficient deterrence from cheating. To implement this idea we have to disregard both expected payoff and imitation considerations in some instances. Thus, this rule stands in a sense even further from rationality than average payoff principle. As we shall see further, the effective punishment rule is also more favorable to compliant behavior. The rule has the same information structure as the previous one.

As a result, the law of motion for q is given by

$$q_{\tau+1} = q_\tau \left(1 - p_\tau\right) \left(3 \left(1 - q_\tau\right)^2 + 3 \left(1 - q_\tau\right) \left(1 - p_\tau\right) q_\tau + \left(1 - p_\tau\right)^2 q_\tau^2\right). \tag{13}$$

This is aggregate population dynamics, and it can be shown that this rule is stabilizing as well. The interior steady state is

$$\bar{q}\left(p\right) = \frac{3(1 - p^2) - \sqrt{1 + 12p - 18p^2 + 8p^3 - 3p^4}}{2(1 - p^3)}. \tag{14}$$

In the **commitment** case the long run noncompliance function is

$$\hat{q}\left(p\right) = \begin{cases} \bar{q}, & p \in \left[0, \frac{2}{3}\right); \\ 0, & p \in \left[\frac{2}{3}, 1\right]. \end{cases}$$

There is no corner solution in this case, $p^* \in \left[0, \frac{2}{3}\right]$. The effective punishment rule is contributing more to the honest reporting, and it is of no surprise that the optimal probability of auditing is lower here for the same parameter values.

The payoffs of the tax authority for the effective punishment are increasing in the magnitude of fine slower than for the best average. As a result, the interval of s for which the tax revenue of static game exceeds that of dynamic is larger for the effective punishment rule, holding all the parameters constant.

Convergence features are not altered either: to reach no cheating state from the middle takes 594 periods now (compared with 597 before); to get to all cheating takes 3 periods (14 before). The latter, however, can not be compared directly, as for the best average all cheating was attainable at $p \in \left[0, \frac{1}{\sqrt{3}}\right]$ and computed for $p = \frac{1}{\sqrt{3}}$; for the present rule it can only happen for $p = 0$.

In the **no commitment** case the system converges (in continuous time) to the steady state with lower probability of auditing than with the previous rule. p^{ss} comes from as a solution of the third-order polynomial $\hat{q}(p) = \mu_1$. From the phase portrait on Figure 2 it is clear that this value is lower than for the best average principle.

3.3   Endogenous switching between rules

The learning rules implicitly reflect the attitude of the taxpayers towards compliance. So far we were treating these rules as given. A more realistic assumption, however, would allow for increased tolerance towards cheating when cheating is widespread. Consider the simplest case, when the average payoff principle is used in case more than a share $a \in (0, 1)$ of population is cheating, and effective punishment rule is used otherwise[9].

Interestingly, the dynamics remains qualitatively unchanged even for this endogenously mixed rule. For high values of a, that is when switching between rules happens with not too high cheating, the stable set reminds the one for effective punishment principle. Correspondingly, for low values of a the dynamics is similar to the one generated by average payoff principle. Despite some irregularities that may arise in the behavior of the system, the general cycling pattern remains intact, that can best

---

[9]Recall that in previous sections we have shown that average payoff rule reflects better attitude towards cheating, and the effective punishment reflects better attitude towards compliance.

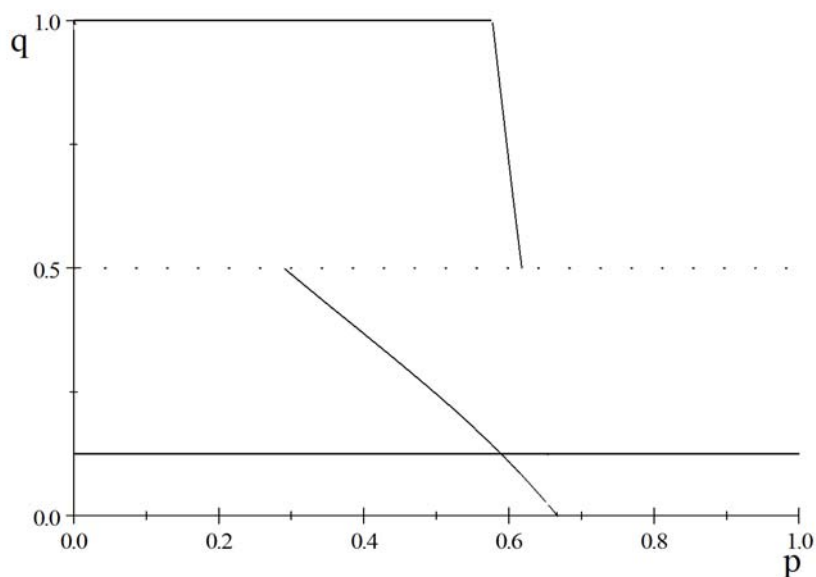be seen on the phase diagram below, $a = 0.5$:

Figure 3. Endogenous switching between average payoff and effective punishment

The problem of committed authority becomes a bit more complicated, as now the payoff is a discontinuous function of auditing probability. Nevertheless, proposition 2 remains valid as long as $\bar{q} \neq a$. An interesting task related to endogenous switching would be to analyze a decision of welfare maximizing government that can induce change in the attitude towards compliance at certain cost. This is however beyond the scope of this paper, and we leave it for future research.

3.4   Meeting m others: Popularity principle

When $m + 1$ people meet (and m is substantially larger than 2), we can specify an imitation rule that requires minimal information about the individual, and namely only whether he/she was caught cheating. Assume that the availability of this information is assured by the tax authority for the purpose of deterring the others. This seems plausible, as the big evasion scandals are normally exposed to the light of public attention[10]. Since this rule does not require the knowledge of the tax report, it can cover groups of substantial size, e.g. colleagues at a single firm or even the members of the same profession. Then the ease of information dissipation will be

---

[10]For ample evidence type e.g. "tax evasion" in the search engine of *Financial Times* at www.ft.com

reflected in the group size in our model, and we can see how it influences cheating and auditing.

The intuition behind the rule is straightforward: an agent gets scared and chooses honesty, if she is caught, or if she observes more than k^* of other caught agents. The rule is then the following. For a not caught taxpayer: if more than k^* caught individuals are observed, play honest in the next round, if less or equal - play cheat; for a caught taxpayer: play honest.

The probability to observe less or k^* caught individuals is defined by

$$\Pr(k \leq k^*) = \sum_{i=0}^{k^*} \binom{m}{i} (pq)^i (1-pq)^{m-i}.$$

Then cheating is evolving according to

$$q_{\tau+1} = (1 - q_\tau p_\tau) \Pr(k \leq k^*). \tag{15}$$

Though the rule is stabilizing, the problem with its dynamics is that once the system comes close to extreme values of q (0 or 1), it is jumping between "almost all cheating" and "almost all honest" states in every period. This problem obviously states from an 'epidemic' nature of the specified principle: once there are very many cheaters, almost everybody meets a caught cheater, and then all those switch to playing honest. But once almost everybody is playing honest, almost nobody meets a caught cheater, and then almost everybody is playing cheat. Proposition 3 still applies, as continuous time serves as a natural stabilizer in this case.

The usual method to make the dynamics more smooth is to introduce some kind of inertia into the system, just like it was already done from the side of the tax authority. So, let us say that with probability β every unpunished individual changes his/her strategy according to already specified rule, and, correspondingly, with probability $1 - \beta$ plays the same strategy as in the previous period. As before, punished people switch to compliance with probability 1, and thus do not exhibit any inertia.

Then in every period $(1 - \beta)(1 - p)q + \beta q(1 - p) \Pr(k \leq k^*)$ cheaters remain cheaters plus $\beta(1 - q) \Pr(k \leq k^*)$ honest people switch to cheating. The dynamics is described by

$$q_{\tau+1} = q_\tau(1 - p_\tau)(1 - \beta) + \beta(1 - q_\tau p_\tau) \Pr(k \leq k^*). \tag{16}$$

For small enough values of β it converges to a steady state (cycle in discrete time) rather than jumps between two extreme values. For simplicity we further consider the

case when observing one caught individual is enough to deter from evasion ($k^* = 0$). The dynamics is then

$$q_{\tau+1} = q_\tau(1 - p_\tau)(1 - \beta) + \beta(1 - q_\tau p_\tau)^{m+1} \tag{17}$$

In **no commitment** case we again observe small cycles around the steady state. Compared to the previous imitation rules, the steady cheating line is shifted to low cheating - low auditing corner, meaning that steady state is more likely to have low probability of auditing. This comes from two factors: inertia in decision making β and number of people to meet m. Notice, however, that even for $p \to 1$ cheating is not eliminated completely. Indeed, for $p_\tau = 1$ $q_{\tau+1} = \beta(1 - q_\tau)^{m+1}$, so that $q = 0$ only for $\beta = 0$, which is impossible. Hence, for large auditing probabilities m-rule results in larger cheating than 3-rules. This seemingly strange result stems from poor information the individuals possess: if nobody is cheating, nobody is caught, so in the next period β of individuals will cheat.

The steady state auditing is decreasing in m, and hence in the speed of information dissipation. Indeed, from (17) it can be computed that $\frac{dp^{ss}}{dm} < 0$[11].

In the **commitment** case, then, cheating can be decreasing or increasing over time depending on whether $q_\tau > \hat{q}(p^*)$ or the opposite. $\frac{dp^*}{d\beta}$ has an arbitrary sign. Since honest reporting is favored more by m-rule, we expect optimal auditing to be lower for the same parameters. The magnitude of fine is almost irrelevant under the present imitation rule, since there is no information about payoffs, and people are deterred from evasion by observing caught cheaters regardless of financial costs of being caught. Increasing the number of people met in this rule also brings about less cheating, because seeing more people means higher chance of observing a caught one. Formally, $\frac{dp^*}{dm} > 0$.

3.5 Parameterization

In the following we parametrize our model to study the dynamics quantitatively: $s = 0.8$, $t = 0.3$, $\frac{c}{H-L} = 0.06$ ($\mu = 0.2$), $\gamma = 0.5$. The fine is usually up to the amount of tax evaded, and I take 20% less than the whole. The income tax rate ranges from 0.1 to 0.5 across developed countries; the measures for both $\frac{c}{H-L}$ and γ are bound to be arbitrary, since in reality the auditing function depends on many more variables

---

[11] $\frac{dp^{ss}}{dm} = \frac{\beta(1-\mu_1 p^{ss})^{m+1}\ln(1-\mu_1 p^{ss})}{1-\beta+\beta(m+1)(1-\mu_1 p^{ss})^m \mu_1}$

than just income, and there is a continuum of income levels rather than two. A convenient way to think of the first measure is as of what share of audited income has to be foregone for the auditing itself. Andreoni et al. (1998, p. 834) take 0.05 as an example, I think of 0.01 to 0.1 as a possible range. Finally, γ to certain extent reflects the income distribution, and 0.5 gives an extreme case where there is an equal number of the rich and the poor.

For the average payoff principle then $p^* = 0.62$. Auditing is increasing in the cost - tax bill ratio, $\frac{dp^*}{d\mu} = 0.15$ (0.043 for the effective punishment), and decreasing in the amount of fine, $\frac{dp^*}{ds} = -0.18$ (-0.258 for the effective punishment). Intuitively, faced with higher fine or lower auditing costs, the taxpayers will cheat less in steady state, hence there is no need for the tax authority to commit to a higher auditing probability.

For the parameter values chosen, the tax authority is better-off with imitating taxpayers for the magnitude of fine smaller than 0.5 and worse off for the magnitude larger than 0.5. This is quite intuitive, since low (high) values of s result in large (small) auditing probability of static no cheating equilibrium; auditing, in turn, is costly to implement. In dynamic setting the auditing probability for given parameter values hits the upper bound of $\frac{2}{3}$, and hence is independent of the surcharge rate, except for the values of s close to 1. Consequently, "static" revenue is increasing with the fine, whereas the "dynamic" is staying constant.

For the high values of γ the picture remains the same, except that now for very large values of fine the "dynamic" revenue rises so much that it exceeds the "static" one. Finally, with decrease in μ the solution with p strictly less than $\frac{2}{3}$ is obtained for larger and larger set of s values, approaching $s \in (\frac{1}{2}, 1]$. Correspondingly, the superiority of "static" revenue is preserved only at $s = \frac{1}{2}$ in the limit (μ close to 0).

In the no commitment case ($\mu_1 = 0.125$) the steady state value of p is equal to 0.625 for the effective punishment rule compared with 0.659 for the average payoff rule. The payoff of tax authority is slightly higher when it is able to commit[12].

The discrete time dynamics is presented for two rules are presented on the Figures

---

[12]Generally for the interior, the condition for superiority of commitment for tax authority is $(1 + s - \mu) \left(\bar{q} p^* - \mu_1 \bar{q}^{-1} (\mu_1) \right) > \mu \left(p^* - \bar{q}^{-1} (\mu_1) \right) \frac{1-\gamma}{\gamma} + \bar{q} - \mu_1$.
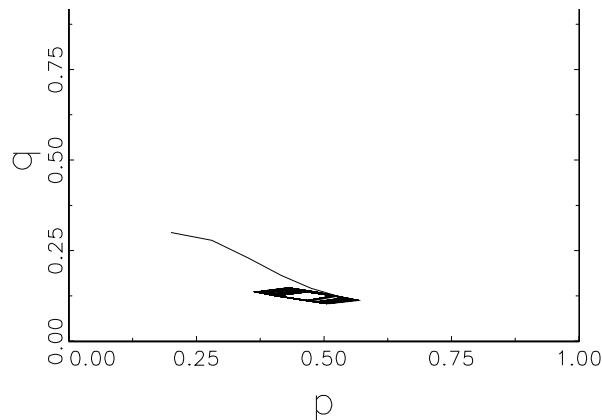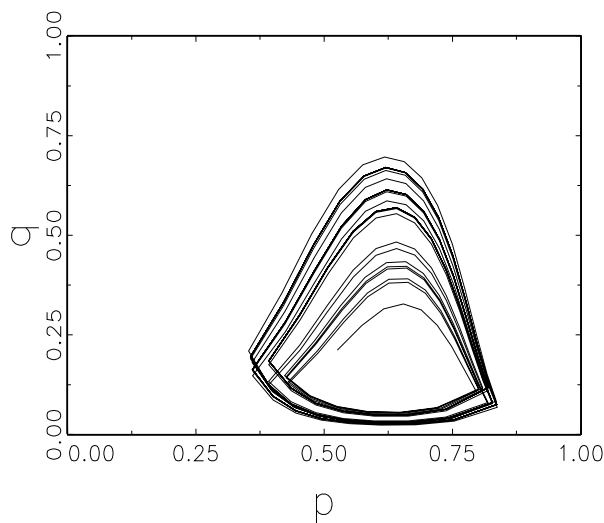
2, 3.

Figure 4. Dynamics by average payoff principle. Figure 5. Dynamics by effective punishment principle.

For the m rule p^* is increasing in β: tax authority has to audit more, if larger part of individuals is reconsidering their decision at every period. The change of surcharge rate is not changing p^*, and the equilibrium auditing is lower than before: 0.43 compared with 0.65 for the first rule (the difference between rules is increasing with the cost of auditing c). Note that this stems mostly from higher number of people who interact, rather than from the different information structure of the rule. Indeed, for $m = 2$ optimal probability is 0.61, not substantially lower than for the other rule.

4 Generalized problem

The analysis presented above is by no means confined to tax evasion. A general compliance problem can be addressed within the same framework: the crucial ingredients are a single monitoring (controlling) authority and a large population of interacting agents that have an option to comply. These can be criminals vs. police, corrupted officials vs. anti-corruption body, firms not complying with quality or safety regulations vs. corresponding monitoring authorities, polluting producers vs. environmental authority, traffic violators vs. road police, free riders vs. controllers in public transportation. The generalized problem can be then formulated as follows.

In period 0 the otherwise homogenous population of agents is playing comply or not comply. The shares of non-compliant q_0 and audited p_0 agents are given

24

exogenously. They determine period 0 payoffs of agents and the authority completely. Between period 0 and period 1 the authority is updating its belief about the share of compliant agents, the agents meet in groups of n and receive information specified by a given behavioral rule. In period 1 the authority performs auditing that either maximizes its expected payoff in this period (no commitment) or maximizes its long-run payoff (commitment). The agents choose comply or not comply according to the behavioral rule. The game is then repeated infinitely.

In the present formulation the agents do not submit a report and are not divided in income groups. Each may choose noncompliance, however, the authority has to *prove* the fact of non-compliance, even if it *knows* that it takes place. Such setup is especially appealing in application to corruption or free riding: a bureaucrat has to be caught receiving a bribe in order to be penalized, even if it is a common knowledge that she/he is corrupted; a passenger without a ticket will not be charged a fine unless she meets a controller.

Propositions 1 and 3 go through in the generalized setup, as the payoffs play only an implicit role in them. The dynamics in case of commitment remains therefore the same. Proposition 2 as well as assumption 2 have to be changed, as the payoff of the authority in generalized form is

$$(1 - \hat{q}(p))V + p\hat{q}(p)F - cp \tag{18}$$

rather than (**??**). Here $V > 0$ is the value of compliance for the authority, and $F > 0$ is the value from detection and punishment. The value of noncompliance is normalized to zero.

The first order condition is

$$(p^*\bar{q}'(p^*) + \bar{q}(p^*))\, F = \bar{q}'(p^*)V + c, \tag{19}$$

and we can observe the equality of generalized marginal benefit and marginal costs. The second order condition is

$$d := (p^*\bar{q}''(p^*) + 2\bar{q}'(p^*))\, F - \bar{q}''(p^*)V < 0. \tag{20}$$

The comparative statics is

$$\frac{dp^*}{dc} = \frac{1}{d} < 0, \tag{21}$$

$$\frac{dp^*}{dF} = \frac{\bar{q}'V'(F) - (p^*\bar{q}' + \bar{q})}{d}; \tag{22}$$

25

that is higher auditing costs will always have higher direct effect of reducing monitoring in this setup. The effect of fine is the same as in proposition 2, if V and F are independent. If they are positively related, the impact of the fine is increased (becomes more positive or less negative); in case of negative relation the impact is decreased (becomes less positive or more negative).

In case of no commitment, the switch in best response happens at

$$\hat{\mu}_1 = \frac{c}{F}. \tag{23}$$

The steady state is again then described by (10) and corresponding dynamics is qualitatively the same as in tax evasion problem.

5   Conclusion

The model presented in the paper is designed to capture a number of features of reality, which were largely neglected in the literature on tax evasion, and especially in the game-theoretic approach to the problem. These features are social interaction, poor knowledge of auditing probability, asymmetry in the behavior of two parties under consideration, and intertemporal nature of the tax evasion decision. The interaction in the model is learning each others' strategies and payoffs. This allows individuals to make decisions without acquiring information about auditing probability. Moreover, with simple imitation rules specified in the game, people also avoid costs of processing information, as they effectively know what decision to take without solving complicated maximization problems.

The model may rationalize decrease of auditing probability (Slemrod 2007) and compliance (Graetz, Reinganum and Wilde 1986) observed in the US over past decades as out-of-equilibrium dynamics. The model can also potentially explain "too little" cheating by taxpayers: having initially overestimated auditing probability, they "undercheat" for a long time due to the inertia and imperfections of the learning rules. All these results hold for different specifications of the learning rule (our rules differ in how much people are afraid of being caught and how much information they can learn from each other).

When we allow the tax agency to commit to a certain probability of auditing, positive cheating may arise in equilibrium. This seems more plausible than the result obtained in the most of static commitment models. Such models usually have zero

26

cheating of audited taxpayers in equilibrium. Moreover, as opposed to the models in the literature, the comparative statics with respect to tax rate does not contradict empirical evidence (cheating is increasing with tax) for a set of parameter values.

However, the relevance of the model to the policy has its obvious limitations. For instance, nothing can be said about the extent of inertia in auditing decision, though this could probably be empirically testable. Without good feeling about the inertia parameter and the learning rule we can not say much about the precise form the dynamics takes. Furthermore, in reality there are different groups of taxpayers and different audit classes. Whereas our analysis can be applied to each group separately, it does not take into account possible inter-group interactions.

In general, the dynamic approach to tax compliance games reopens a whole bunch of policy issues. Are the recommendations of equilibrium theory valid, if the systems does not converge to an equilibrium? Are some changes in the existing taxation worth undertaking, if we take into consideration not only difference in benefits between initial and final states, but also the costs of transition? Can the decision rules of the tax authorities and the learning mechanisms governing taxpayers behavior be manipulated in the way to achieve maximal social welfare?

As a building block for more general models, the behavioral approach can be employed in the studies on how the government can ensure higher degree of trust in society (and less evasion as a result), how it can provide optimal (from the point of view of social welfare) level of public goods, how it can bring about faster growth of an economy. For this it would be necessary to consider more complicated government (and hence tax authorities) strategies, involving more than one period memory, and possibly heterogenous taxpayers.

Finally, the approach taken by no means limits us to consideration of tax evasion. The model can in principle be applied to violations of law other than tax evasion, if they have properties of unobservability, costly monitoring and interaction of agents. The examples here include crime, corruption, employment and environmental protection, traffic rules violation, financing of public utilities.

References

[1] Allingham M., Sandmo A. Income tax evasion: a theoretical analysis. Journal of Public Economics, November 1972, pp. 323-338.

[2] Andreoni J., Erard J. and Feinstein J. Tax Compliance. Journal of Economic Literature, June 1998, pp. 818-860.

[3] Bardsley P. Tax compliance games with imperfect auditing. The university of Melbourne Department of economics research paper no.548. January 1997.

[4] Clotfelter C. Tax Evasion and Tax Rates: an Analysis of Individual Returns. Review of Economics and Statistics, 65(3), pp. 363-73, 1983.

[5] Cowell F. The Economics of Tax Evasion, MIT Press, 1990.

[6] Cressman R., Morrison W. and Wen J-F. On the Evolutionary Dynamics of Crime. Canadian Journal of Economics, vol. 31, no. 5, November 1998.

[7] Dubin J., Graetz M. and Wilde L. The Effect of Audit Rates on the Federal Individual Income Tax, 1977-1986. National Tax Journal, 43(4), pp. 395-409, 1990.

[8] Engel E., Hines J.. Understanding Tax Evasion Dynamics. NBER Working Paper No. 6903, January 1999.

[9] Erard B., Feinstein J. Honesty and Evasion in the Tax Compliance Game. Rand Journal of Economics, 25(1), pp. 1-19, 1994.

[10] Fudenberg D., Levine D. The theory of learning in games. MIT Press, 1998.

[11] Giles D., Caragata P. The learning path of the hidden economy: the tax burden and the tax evasion in New Zealand. Econometrics Working Paper 9904, Department of economics, University of Victoria, Canada.

[12] Graetz M., Reinganum J. and Wilde L. The Tax Compliance Game: Towards an Interactive Theory of Law Enforcement. Journal of Law, Economics and Organization, 2(1), pp. 1-32, 1986.

[13] Hirsch M., Smale S. Differential equations, dynamical systems, and linear algebra. New York : Academic Press, 1974.

[14] Manski C. Economic Analysis of Social Interactions. Journal of Economic Prospectives, 14(3), pp. 115-136, Summer 2000.

[15] Poterba J. Tax Evasion and Capital Gains Taxation. American Economic Review, Papers and Proceedings, 77, pp. 234-39, 1987.

[16] Reinganum J., Wilde L. Equilibrium Verification and Reporting Policies in a Model of Tax Compliance. International Economic Review, 27(3), pp. 739-60, 1986.

[17] Sánchez I., Sobel J. Hierarchical Design and Enforcement of Income Tax Policies. Journal of Public Economics, 50(3), pp. 345-69, 1993.

[18] Schlag K. Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-armed Bandits. Journal of Economic Theory 78, pp.130-56, 1998.

[19] Slemrod J. The Economics of Corporate Tax Selfishness. National Tax Journal 57(4), December 2004.

[20] Slemrod J. Cheating Ourselves: The Economics of Tax Evasion. Journal of Economic Perspectives, Vol. 21, No. 1, Winter 2007

[21] Slemrod J., Yitzhaki S. Tax Avoidance, Evasion and Administration. Handbook of Public Economics, volume 3, 1423-1470, North Holland, 2002.

[22] Schneider F., Enste D. Shadow Economies: Size, Causes, and Concequences. Journal of Economic Literature, pp.77-114. March 2000.

[23] Wane W. Tax evasion, corruption, and the remuneration of heterogeneous inspectors. World Bank Working Paper 2394, 2000.

[24] Witte A., Woodbury D. The Effect of Tax Laws and Tax Administration on Tax Compliance. National Tax Journal, 38, pp.1-13, 1985.

6 Appendix - Proofs

Lemma 1. Under assumption 1, for stabilizing rules $f_p(q,p) < 0$; for destabilizing rules $f_p(q,p) > 0$.

Proof. By implicit function theorem, $f_p(\bar{q}(p),p) = -\bar{q}'(p) f_q(\bar{q}(p),p)$. For stabilizing rules then $f_p(\bar{q}(p),p) < 0$. Suppose $\exists q | f_p(q,p) < 0$. By continuity (a derivative of a polynomial is again a polynomial and hence continuous), $\exists q_0 \in (q, \bar{q}(p)) | f_p(q_0,p) = 0$, which contradicts the uniqueness of the

steady state. This proves the lemma for $p \in (\bar{q}^{-1}(1), \bar{q}^{-1}(0))$. For $p \in [0, \bar{q}^{-1}(1)]$[13], by continuity in both arguments $\exists \delta, \varepsilon > 0 | f_p(1-\delta, \bar{q}^{-1}(1)-\varepsilon) < 0$. Suppose $\exists p, q | f_p(q, p) > 0$, then again by continuity $\exists q_0 \in (q, 1-\delta), p_0 \in (p, \bar{q}^{-1}(1) - \varepsilon) | f_p(q_0, p_0) = 0$. But this would mean that q_0, p_0 constitute a steady state, that contradicts the corner solution. The proof for $p \in [\bar{q}^{-1}(0), 1]$ is completely analogous. For destabilizing rules $f_p(\bar{q}(p), p) > 0$ and the rest of the proof follows the same logic.

Proof of Proposition 3. To investigate stability of the steady state analytically, we have to make two approximations. First, consider the system in continuous time: this makes sense, if we imagine that both the tax authority and individuals update their evasion or auditing decisions every day, rather than fixing it once for a whole year. We can rewrite our system of equations as

$$q_{\tau+\Delta} = q_\tau + \Delta f(q_\tau, p_\tau),$$
$$p_{\tau+\Delta} = p_\tau + \Delta g(q_\tau, p_\tau);$$

and letting Δ be very small ($\frac{1}{365}$, if we think of daily updating), in the limit we obtain

$$\dot{q} = f(q, p),$$
$$\dot{p} = g(q, p);$$

where $f(q, p)$ is defined by the learning rule and $g(q, p) = \alpha(BR(q) - p)$.

The stability matrix of this system is $\begin{pmatrix} \frac{\partial \dot{q}}{\partial q} & \frac{\partial \dot{q}}{\partial p} \\ \frac{\partial \dot{p}}{\partial q} & \frac{\partial \dot{p}}{\partial p} \end{pmatrix} = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} = \begin{pmatrix} f_q(q, p) & f_p(q, p) \\ \alpha BR'(q) & -\alpha \end{pmatrix}$.

The problem with this formulation is that the best response function is not continuous at the point of steady state, so we can not compute $BR'(q_{ss})$. To go around it, we can make the second approximation: instead of the discontinuous best response we take a continuous function $ABR(q) = \Phi\left(\frac{\bar{c}(q)-c}{\sigma}\right)$, which approaches $BR(q) = \begin{cases} 0, & \text{if } \bar{c} < c \\ 1, & \text{if } \bar{c} > c \end{cases}$ with $\sigma \to 0$. Conventionally, Φ is cumulative distribution function of a standard normal random variable. Then $ABR'(q) = \phi\left(\frac{\bar{c}(q)-c}{\sigma}\right)\frac{\bar{c}'(q)}{\sigma}$. Recalling the expression for $\bar{c}(q)$ and evaluating at steady state ($\bar{c}(q_{ss}) = c$), we get

---

[13]In case $[0, \bar{q}^{-1}(1)]$ is not an empty set.

$$ABR'(q_{ss}) = \phi(0)\frac{(1-\gamma)c}{\sigma q(1-\gamma+q\gamma)} \Rightarrow a_{21} \approx \frac{\alpha(1-\gamma)c}{\sqrt{2\pi}\sigma^2 q(1-\gamma+q\gamma)}.$$

Note that we can make a_{21} (since it is positive) arbitrary large by making σ small enough and thus getting better approximation of initial best response function.

Now we are ready to address the question of stability of the steady state. If the real parts of both eigenvalues of the stability matrix are negative, the steady state is stable (see, for example, Hirsch and Smale (1974). The eigenvalues of our system are

$$\lambda_{1,2} = \frac{1}{2}\left(a_{11} + a_{22} \pm \sqrt{(a_{11}-a_{22})^2 + 4a_{12}a_{21}}\right).$$

Note that by Lemma 1 in steady state $\frac{a_{12}}{a_{11}} > 0$. Then, quite intuitively, the stabilizing learning rules ($f_q(q,p) < 0$) will lead to convergence. Indeed, for such rules $a_{11} + a_{22} < 0$ and $(a_{11}-a_{22})^2 + 4a_{12}a_{21} < 0$. Hence, both eigenvalues have negative real parts[14] - our steady state is stable in continuous time. For the rules that are destabilizing ($f_q(q,p) > 0$), we shall have no convergence. Indeed, in this case one eigenvalue is positive, the other is negative - the linarized system is a saddle. Note that our results hold for any $\alpha \in (0,1)$.

Condition for interior solution - average payoff rule In order to get interior solution, we have to get more tax revenue there than at $p = 0$. That is, the following inequality should hold:

$$\begin{aligned} \gamma(1-q)tH + p(q\gamma(tH + st(H-L)) + (1-\gamma)tL) \\ -c(p(q\gamma+1-\gamma)) + (1-p)(q\gamma+1-\gamma)tL > tL \end{aligned} \quad (24)$$

Collecting the terms and using the definition (4) for μ we arrive at

$$\frac{1-q+pq(1+s)}{p(\frac{1}{\gamma}-1+q\gamma)} > \mu, \quad (25)$$

which after substituting q with its steady state value $\bar{q}$ from (12) becomes

$$\frac{-1/p + (2-3p)s + 2}{1 - 3p^2 + \frac{1}{\gamma}(1-3p+3p^2)} > \mu.$$

As $p^* \in \left(\frac{1}{\sqrt{3}}, \frac{2}{3}\right)$, we can check the inequality at the ends of the interval to obtain

$$\mu < \gamma(s+1), \mu < 3\frac{\gamma}{2(1-\gamma)}.$$

Both conditions show that a corner solution obtains for large μ and small γ.

---

[14]We do not make any statements about the corner solutions at this point.

Convex auditing function. With a convex auditing function $C(.)$ the myopic best response of tax authority is no longer jumping, but is a smooth function of q:

$$BR(q) = \frac{1}{q\gamma + 1 - \gamma} C'^{-1}\left(\frac{q\gamma(1+s)t(H-L)}{q\gamma + 1 - \gamma}\right).$$

and all our insights are preserved, if the slope of the steady auditing line is not too negative:

$$\frac{dq}{dp^{BR}} = \frac{C''(.)(q\gamma + 1 - \gamma)^2}{\gamma((1+s)t(H-L) - C'(.) - C''(.)p^{BR}(q\gamma + 1 - \gamma))},$$

that is if the auditing is not too convex.

7 Appendix - Learning rules

Here we consider two more examples of the learning rules as well as a generic deterministic behavioral rule.

7.1 Proportional imitation

To apply proportional imitation rule (PIR), which was proposed and shown to be optimal by Schlag (1998), we have to modify our setup slightly. The problem with this rule is the need to know the highest and the lowest payoffs of the agents, which was assumed away so far. This does not seem to be a very strong assumption to make: people may know the payoff and still not pursue certain strategy, just because they do not know how to do it.

According to the rule, each agent meets only one other and imitate its strategy with probability proportionate to the payoff difference, if this other performed better. Recall that three payoffs of our game are $(1-t)H - st(H-L)$ if caught, $(1-t)H$ if honest, and $H - tL$ if not caught. The difference between the highest and the lowest is $(1+s)t(H-L)$, evaded tax plus a fine. Then not caught cheater never switches; honest taxpayer meeting not caught one switches with probability $\frac{1}{1+s}$; a caught cheater meeting an honest agent switches with probability $\frac{s}{1+s}$.

Then the law of motion for q is given by

$1 - q_+ = (1-q)[1 - q(1-p)\frac{1}{1+s} + qp\frac{s}{1+s}]$,

where the right hand side is again at time τ.

From this expression, the proportion of cheaters increases, if $p < \frac{1}{1+s}$, and decreases otherwise. Thus, we get the circling around $p = \frac{1}{1+s}$ and $q = q(c)$ again. Interestingly, only with proportional imitation rule the interior rest point is precisely the Nash equilibrium of the static game. It happens because in the present specification the agents possess more information (about payoffs) and have rather sophisticated learning technique.

7.2 Best average with only two people meeting

With only two meeting, we still have at time τ the following types of high income taxpayers: (i) honest, comprising proportion $1 - q_\tau$ of population and receiving payoff $(1-t)H$; (ii) caught cheating, $q_\tau p_\tau$ of population with payoff $(1-t)H - st(H - L)$; (iii) not caught cheating, $q_\tau(1 - p_\tau)$ of population with payoff $H - tL$. We use the best average principle, which in our case actually turns into best payoff (no need to take average). So, there are only following six pairs with corresponding probabilities and outcomes of interaction:

two honest	$(1 - q)^2$	honest
two caught	$(pq)^2$	honest
two not caught	$(1 - p)^2 q^2$	cheat
honest and caught	$(1 - q)\, pq$	honest
honest and not caught	$(1 - q)\, (1 - p)\, q$	cheat
caught and not caught	$p\, (1 - p)\, q^2$	cheat

Thus, the share of cheaters tomorrow is determined according to the following law:

$$q_+ = (1 - p)\, q$$

Thus, the cheaters tomorrow are only non-caught cheaters today. We see honesty prevailing: since p is bounded below by zero, the share of cheaters is non-increasing function of time. Thus, in commitment case we will always get zero cheating in the limit, whereas in no commitment case the cheating will converge to a positive constant value. We conclude that the case of meeting of one other person is a degenerate form of the process under investigation, and it can not capture the features of the evasion problem we are interested in.

7.3 Arbitrary rule with n people in a match

Consider a general deterministic n persons rule without eigen bias (the rule is invariant to the distribution of types in a match). All individuals that belong to the same match exhibit the same behavior. Assume also that we have m different observable types of individuals (in the examples considered we have $m \leq 3$). Call the number of combinations of the types in a match $M := m^{n-1} + \binom{n}{m}$. It can be shown that 2^M distinct rules of this sort can be formulated, with the dynamics represented by $q_{t+1} = \sum_{i=1}^{M} \mathbb{I}_i Q_t(i)$. Here $Q(i)$ is the probability that the combination i of the types occurs, and the particular rule specifies a sequence $(\mathbb{I}_i)_{i=1}^{M}, \mathbb{I}_i \in \{0,1\}$. For a rule not to be degenerate (jumping to a corner immediately), we must have

$$\exists i, \mathbb{I}_i \neq 0; \exists j, \mathbb{I}_j \neq 1.$$

Let us concentrate on the endogenous types, that is audited honest, not audited honest, audited non-compliant and non-audited non-compliant. A period-to-period dynamics resulting from any rule can be generally written as

$$
\begin{aligned}
q_{t+1} &= \sum_{j=0}^{n} \sum_{k=0}^{j} \mathbb{I}_{j,k} n_{j,k} (1-q)^{n-j} q^j (1-p)^{j-k} p^k, \\
f(q,p) &= \sum_{j=0}^{n} \sum_{k=0}^{j} \mathbb{I}_{j,k} n_{j,k} (1-q)^{n-j} q^j (1-p)^{j-k} p^k - q,
\end{aligned}
$$

where $n_{j,k}$ is a fixed coefficient characterizing number of permutations for the same combination of types.

As can be seen, the function $f(q,p)$ is a polynomial of the order at most n in each of the arguments. Adding exogenous types (like rich-poor, self-employed - employees) does not change this result, as they enter the function in a similar fashion with simplex coefficients.