

# The Distributive Basis of Tax Compliance

Pablo Beramendi  
Duke University  
and  
Raymond M. Duch  
Centre for Experimental Social Sciences  
Nuffield College  
University of Oxford  
OX1 1NF Oxford UK  
raymond.duch@nuffield.ox.ac.uk

\*

January 17, 2015

---

\*Paper prepared for the *Symposium in Celebration of Margaret Levi*, University of Washington, Seattle, October 10th, 2014. **Very Preliminary Work**



# 1 Introduction: Why Tax Compliance Matters

The ability of governments to function depends on the resources they have. Without sufficient resources there is no policy choice, and without choice there is no politics. The very development of markets, that is the size of the pie to be allocated via future political struggles, requires effective legal and fiscal institutions (?). Unsurprisingly, the fiscal development of the state becomes in itself a major political battleground for inequality and prosperity around the world are as much about what governments can do as they are about what governments cannot do.

Figure 1 below plots an indicator of tax capacity against the overall tax revenue yield as a share of economy (left panel) and a measure of post tax and transfers income inequality (right panel). Unequivocally, the institutional capacity of the state drives both the level of revenue collection and the ability to combat inequality. Interestingly, Southern Europe provides a clear transition, a missing linkage as it were, between developed democracies and a world in which the governments choice set becomes narrower and narrower.

Figure 1: Inequality and Capacity

### Capacity and Inequality



One can think of compliance as the slope of the fitted line in the left panel: how much the state is able to generate given a fixed level of capacity. A steep line implies that citizens are responsive to state's efforts to increase their fiscal capacity. A flat line implies that citizens are not sensitive to increases in effort to generate resources. Holding states' efforts constant, compliance drives the effective capacity of the state, and with it, much of what we refer to as politics. Without compliance, there are less means for policy, and ultimately less room for using policy to expand the citizens' capability set. This is precisely what makes compliance so central for political economy, yesterday and today.

To remain effective the state rests on the willingness by citizens to observe their obligations towards the rest of the political community. This fundamental intuition was central to Levi's notion of *quasi-voluntary compliance*. Citizens comply not only because of the monitoring ability of the state but also, and primarily, because of a normative endorsement to collective exchanges that they perceive as fair and good for the community. In her own words: "*Citizens comply less when they perceive that others are not paying their fair share and the state is not delivering the perceived goods*" (?). Insofar as this combination of strategic and normative considerations lead a sufficient number of people to comply, the process is reinforced over time and quasi-voluntary compliance becomes self-enforced.

In this paper we revisit the question of the mechanisms driving self-enforcement in quasi-voluntary compliance. We contribute to an extensive literature on an old puzzle in political economy: *why do people pay taxes?* In particular, why do high income citizens pay taxes even in the absence of auditing institutions? What explains the fact that actual tax compliance is higher than the level identified as a fair norm by experimental researchers? In addressing this questions, we propose a *distributive logic* of tax compliance. In essence we argue that beyond arguments about the monitoring ability and the efficiency of state institutions and normative concerns for the worse-off in society, the

distributive incidence of the public sector over the long run drives the level of compliance. In other words, the distribution of net benefits (as determined by the progressivity of taxes and transfers) is a key determinant of what individuals perceive as collectively fair, and therefore drives their willingness to contribute to the common pool. We test the argument on the basis of a series of lab and online experiments in which subjects are randomly assigned to different fiscal treatments, treatments that approximate the variation in the design of national fiscal systems around the world.

By exploring the individual responses in terms of compliance to the specifics in the design of different fiscal systems, the paper makes several contributions. First, it offers a deeper exploration of some of the micro-level assumptions in the macro PE literature on tax capacity and tax compliance. Second, it uses the macro comparative taxation literature to illuminate several important mechanisms driving compliance at the individual level, thereby providing a bridge between behavioral economic approaches and macro-political ones. Third, by performing this blending experimentally, the paper contributes to disentangle important elements of the micro-logic of state capacity building free of the important limitations affecting observational research, most prominently in this case the feedback loops connecting state performance and citizens's dispositions to each other. The rest of the paper is organized as follows. We first present the argument underpinning the *distributive logic of tax compliance* and derive our theoretical expectations. Second, we describe in detail the design in our experiment and outline the main features of both the lab and survey implementations. Third, we present our main findings across both lab and survey analyses. The final section draws the main implications from the results and outlines future lines of research.

## 2 New Motivation Material from Ray

The interesting puzzle is the right hand side of Figure 1 – low tax capacity states have high levels of inequality – high capacity states have low levels of inequality. I think we can make a significant contribution to understanding why this is the case. And I think we should motivate the discussion around this specific puzzle. There are different causal stories suggesting how this state of affairs arises:

One conjecture is that there is a prevailing preference, or maybe tolerance, for high versus low capacity states (i.e., overall tax burdens) – or there is serendipity – somehow a collectivity unfortunately ends up in a high capacity/tax burden state or fortunately ends up in a low capacity/tax burden state.

It then follows that preferences for redistribution are conditioned on state capacity:

- in general, given a low capacity state, the stable equilibrium is one in which self-interested choices favor compliance when redistribution levels are modest – modest redistributive taxation tends to increase compliance – why is this the case? – progressivity tends to lower tax compliance – and we can speculate why this might be the case
- and in general, in a high capacity state, self-interested compliance decisions favor the state in which redistribution is quite aggressive – hence progressivity tends to raise tax compliance – and we can speculate as to why this might be the case

Our interest is teasing out the extent to which this initial condition contributes to the stable equilibrium pictured in Figure 1 . The problem is that in the real world things are more messy – preferences for redistribution and the capacity of the state are not independent – so high inequality and low capacity may jointly be determined by redistributive preferences. The logic goes as follows: the electorate is ambivalent towards redistribution; this is translated into a low capacity state or a state with a low tax burden;

inequality persists without an activist/taxing state.

In order to get some leverage on this puzzle and avoid confounding the influence of capacity and preferences we designed a set of experiments in which subjects are randomly assigned to capacity and redistributive treatments which makes the identification of causal effects much less problematic.

## **2.1 Tax Compliance Under Different Redistributive Regimes: BD Self-interested Model**

When should we expect taxpayers to comply with redistributive taxation and hence when should we expect high levels of redistribution to be sustainable and when should we expect low levels of redistribution to be sustainable? In our model, complying with different levels of redistribution is driven by self-interested assessment of the costs and benefits associated with compliance we argue these are determined by the interaction of capacity and progressivity of redistributive taxation.

Tax payer calculus in this tax compliance game:

- recognize that the equilibrium decision in all of these rounds of the game is to declare no income (I think this is the case even when the audit rate is 10 percent);
- on average this almost never happens positive income is declared on average by all subjects;
- what is interesting here is how this compliance varies across our treatments or why does deviation from equilibrium play vary across treatments;
- self-interest model of compliance with redistributive taxation: this is essentially a one-shot game so what we are likely capturing here is:

- the subjects assessments of the net costs/benefits of compliance if compliance is costly then the subject should cheat
- net costs/benefits vary across the treatments, which correspond to fiscal/benefits regime types under what circumstances is compliance less costly worth the “risk?”
- if subjects respond to these net costs/benefits as hypothesized then the experimental results will lend credence to a self-interested model of support for redistributive preferences
- this provides an understanding of why the right hand graph in Figure 1 represents an equilibrium outcome

### **2.1.1 Low Tax Burden State**

#### **Low Progressivity**

- Rich: rich are more likely to cheat under all scenarios but costs of complying here are low low tax and regressive returns so some might comply
- Poor: compliance to low redistributive taxation is relatively costless here cost of compliance is low and net benefits are low and since dead weight loss is relative small relative to income, net costs of compliance are low even give uncertainty of income distribution, cost of a mistake is quite low I think I am poor but I am actually rich but taxes are low and regressive

#### **High Progressivity**

- Rich: they always cheat costs of complying here are higher if you are rich though so this might bump up their cheating



- Poor: I think this is the way to think about this: if you are poor in a low capacity state, you need to weigh the contribution to your “after tax income of aggressive cheating against what you expect the state can transfer from the rich to the poor under a high progressivity regime – if certain they are poor then cost of compliance is lower but if we add some uncertainty to ones place in income distribution plus some dead weight loss (applied against tax revenues from rich) plus the effects of cheating by the rich, then the expected benefits of redistribution to the poor are not high because taxation is low

### 2.1.2 High Tax Burden

#### Low Progressivity

- Rich: rich are more likely to cheat under all scenarios but costs of complying here are moderate
- Poor: the net benefits of compliance are not that high the state takes a lot of your income (less though because you are poor) and you should expect to get much, if not more, of it back if you are poor although progressivity is low here so transfers are not significant detracting from the benefits though are: 1) dead weight loss which could be significant since the state has taken a lot of your income; 2) compliance to a high tax rate could be costly if one is uncertain of income distribution a mistake could mean a significant transfer of income given high taxes even under low progressivity

#### High Progressivity

- Rich: they always cheat costs of complying here are higher if you are rich though so this might bump up their cheating

- Poor: the net benefits to compliance here are quite significant – the state takes a large share of income (but in a progressive fashion) and redistribution is aggressive so poor benefit on both counts – given that benefits are likely to be high, dead weight loss is less important – but a mistake on assessing ones relative wealth would still be costly

## 2.2 An alternative Rawlsian Model of Tax Compliance

Alternatively tax compliance may simply be conditioned on the progressivity of the redistributive taxes – as the tax rate becomes more progressive and “fair, tax compliance rises

## 3 The Argument: Fiscal Incidence and Tax Compliance

Under what conditions should compliance occur? Taking a step back, a first response emerging from standard political economy models suggests that compliance should *never* be a rational response. Challenging the idea that compliance is, at least partially, in the interest of low income people, a strictly rationalist approach in a one-shot game implies that low income citizens have no incentives to comply. Their optimal strategy is to take the benefits provided by the state without contributing to them. Knowing this, the optimal response by high income citizens is to never comply either. Any one shot game in which contributions to the state are perceived as a zero-sum game would yield such a prediction (e.g. (?)). This benchmark result helps motivate the driving puzzle in this literature as citizens across the world *do* comply. The very notion of quasi-voluntary compliance was a response to this puzzle: Citizens comply more when they fear the disciplinary hand of the state and they perceive the collective exchange with the state as

fair.

Yet what citizens perceive as collectively fair, and therefore, drives them to comply is not obvious. Compliance implies different motives for people at different ends of the income distribution. For low income citizens, compliance is in part self interested as they are net beneficiaries of public goods that otherwise would not be viable; for high income citizens, compliance implies a sacrifice for the sake of the collective. What are the criteria governing such a sacrifice? In many fundamental ways, this is equivalent to asking when the correction of inequalities is justified for those disproportionately bearing its costs.

A first response builds on Rawls' idea of *Justice as Fairness* (??). The notion of justice as fairness is analytically broken down into two principles: the principle of *Fair Equality of Opportunity* and the *Difference Principle*. The former requires that equality of opportunity be effective, not merely formal. The latter establishes, critically, that, provided equality of opportunity exists, the only tolerable inequalities are those that benefit the worse off by eliminating differences that do not emerge from human choices (e.g. natural talent). The combination of these two principles allows Rawls to establish an analytically more precise version of fairness as defined by Marx in the post-socialist society: *From each according to his ability; to each according to his need.*<sup>1</sup> For Rawls, fair implies the exclusive targeting of inequalities that cannot be helped; fair is to focus on those most in need; fair is to correct for the incidence of inequalities that do not derive from conscious human choices assuming a fair distribution of economic opportunity and natural talent. The imprint, however explicit, of this conceptualization of fairness on political economy cannot be overstated. It underlies for instance the inequality aversion parameter in Atkinson's index of Inequality (?), where the parameter  $\alpha$  captures how sensible the index is to the relative welfare of those at the lower part of the income distribution. It underpins Alesina's research and findings on the relationship between fairness and preferences for redistribution. And last but not least, it provides normative

---

<sup>1</sup>Karl Marx, *Critique of the Gotha Program*, 1875.

foundations for targeting benefits to those in need and placing the costs of those benefits disproportionately on the better off, especially if the latter have been benefited by the natural lottery of talent.<sup>2</sup> In other words, fiscal designs with large levels of overall progressivity would reflect, even if loosely (recall that a fair equality is a pre-requisite), a Rawlsian notion of fairness (?). By implication, a more progressive fiscal design should lead to higher levels of tax compliance.<sup>3</sup> The continuous line in Figure ?? below captures the predicted relationship between the progressivity of the fiscal system and the level of compliance under a *Rawlsian* framework.

Implicit to the idea that progressivity is a sacrifice by the better off is the premise that cost and benefits of public interventions are almost perfectly separable. What one class of citizens sacrifices, the other gets. While the rich may derive some utility of their “other regarding” attitudes (?), in material terms, politics, and by implication compliance, remains part of a zero-sum game. In a Rawlsian world, compliance implies a material cost driven by normative considerations. There is however a second, alternative, view of the behavioral basis for tax compliance. Instead of thinking of compliance as a normative deviation from distributive calculations, that is a calculation in which individuals sacrifice utility in the short term for the sake of their principles, we posit in this paper that there in fact exists a *distributive logic* for tax compliance, a logic that does not require the introduction of specific normative assumptions to explain the behavior of middle and high income voters in democracies. This alternative conceptualization of the basis for compliance builds on two assumptions that in turn rest on recent developments in the political economy of taxation and redistribution:

- When making decisions about tax contributions, and compliance is perhaps the

---

<sup>2</sup>For a paper investigating experimentally the different elements of Rawls’ theory and how they combine, see (?). The paper offers partial support to the notion that individuals make distributive allocation decisions on the basis of the maximin principle.

<sup>3</sup>See also ? for an empirical test of the impact of progressivity on tax morale partially consistent with this expectation. They find that progressivity has a positive impact on tax morale and low and medium levels of income.

most important among the latter, individuals care about their relative position in the distribution of expected net benefits in the long run. In other words, life-course considerations play a fundamental role in this individual's utility calculations.

- The state as a public goods provider is not an institutionalized version of Robin Hood, but a risk pooling mechanism over the life-course (?????) to which everyone contributes (according to their ability/means in various degrees) and from which everyone draws (again according to the balance of means and needs in various degrees) at some point over the life course.

Analyzed from these premises, compliance emerges as a rational response to a different set of motives. Rather than a sacrifice for the worse off, compliance works as a long-run investment driven by the possibility that in the medium to long run everyone can be among the *worse off*. Obviously, at any given point in time, not everyone is equally at risk, that is in need to draw from the common pool, but except for a small minority with enough wealth to consistently rely on market options, there will a period in everyone's life in which the tax/benefit<sup>4</sup> balance is clearly positive (this will more the case the higher the life expectancy). In other words, the distribution of net benefits over the long run will be far less skewed than the distribution of net benefits at any given point in time. That the case, middle and high income individuals, which tend to be the ones that discount the future less, have incentives to comply and contribute to the common pool.

There is no gainsaying that significant commitment problems and distributive tensions may undermine these calculations over time. Regarding time inconsistency, self-enforced compliance requires the ability of middle and high income groups to monitor their relative position in the common pool and act politically to adjust it if necessary. Through well established political organizations that bridge interest articulation across legislative terms, they clearly do (?). Regarding distributive tensions, for the common over the life-

---

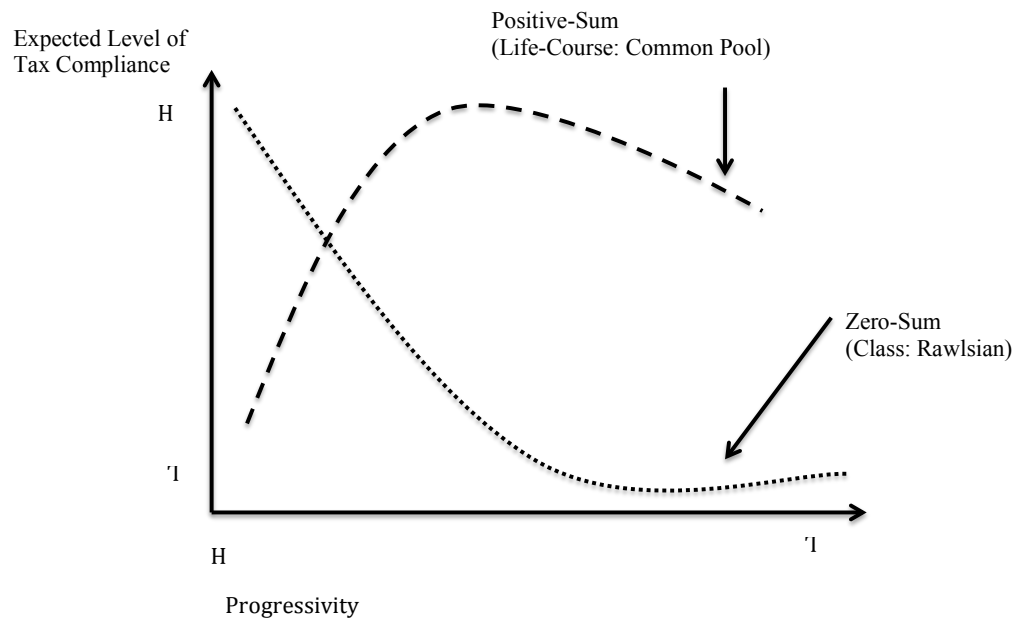
<sup>4</sup>By benefit we refer here to both transfers and in kind services

course to be politically sustainable over time, *everyone* must contribute their fair share, and not just those better off (??).<sup>5</sup> This latter condition, essential for the sustainability of the common pool over time, has direct implications for the link between progressivity and compliance. If the sustainability of PPGG provision requires everyone to contribute, progressivity must be sacrificed, at least in part. Therefore, to the extent that compliance reflects more an investment over the life course than a Rawlsian concern for the worse off, the expected relationship between progressivity and compliance reverses. More progressivity reduces the feasibility of the commons, the expected long run net benefits of high income citizens, and with them the overall level of compliance. In contrast, compliance should peak at medium levels of progressivity, that is to say in systems where both the worse off and middle-to-high income citizens can expect a positive net benefit over the long-run. Finally, in the extreme case of total absence of fiscal progressivity, the overall levels of compliance should decline again: the well to do citizens will comply if, regrettably, they are the main recipients of states' policies, but low income citizens, which in regressive systems are as likely to be taxed, will find no incentives to comply. The dashed line in Figure ?? summarizes the relationship between progressivity and compliance under a *common-pool* framework.

---

<sup>5</sup>A second, increasingly prominent issue, concerns the composition of the population of those with access to the common pool. The increasing importance of immigration as a second dimension fragmenting welfare state support in places like Sweden provides a good illustration of the way high income voters react to broadening the pool to sub sectors of the population that (i) do not belong to the original community, and (ii) have not contributed sufficiently before drawing benefits.

Figure 2: Argument



To summarize, the contrast between the *Rawlsian* and the *Common Pool* approach to compliance yields the following working hypotheses:

1. If the dominant set of motives behind tax compliance reflects a concern primarily for those worse off in society, it should hold empirically that
  - (a) High income people are more likely to comply in more progressive systems
  - (b) As a result, the more progressive the overall fiscal system, the higher the level of compliance in society
  
2. If, on the other hand, compliance reflects primarily long run self-interest in a large common pool of risks, it should hold empirically that
  - (a) High income people are more likely to comply at intermediate levels of progressivity
  - (b) As a result, high levels of progressivity lead to lower levels of compliance in society
  
3. Regardless of the direction specified by each theoretical account, the overall progressivity of the fiscal system (i.e both taxes and transfers) is an important determinant of the levels of tax compliance.

### 3.1 Empirical Strategy

To test these hypotheses we resort to a series of lab and online experiments in which subjects are randomly assigned to different treatments that capture differences in the core dimensions of fiscal systems. The key dimensions over which systems vary are as follows:

- *Level (burden)*: defined by the average tax rate in the system.



- *Progressivity of Taxes*: defined by the elasticity of the tax rate to the level of pre-tax income. In progressive designs, the higher the income, the higher the revenues collected. In other words, the revenue inflow increases more than proportionally with respects to increases in income.
- *Progressivity of Benefits*: defined by the elasticity of the allocation of revenues collected to pre-tax income. In progressive designs, the lower the income the higher the share of benefits received.

Variation in levels is inserted in the treatment as a control for the potential role of the size of the burden on citizens' disposition towards tax compliance. Moreover, level and progressivity are far from independent from each other. It is a well known empirical results that there is a trade-off between the overall size of the fiscal budget and the level of progressivity. In light of these contributions, we know that a state organization that combines a very high tax burden and very high levels of fiscal progressivity is yet to emerge. The reasons behind this *impossibility* are well known and bring us back to the set of necessary political conditions for the state's common pool to remain sustainable over time (??). When defining the different treatments via combinations among these three dimensions, we limit ourselves to a set of realistic alternatives, defined as follows<sup>6</sup>:

- Treatment 1: low burden, low progressivity
- Treatment 2: low burden, high progressivity
- Treatment 3: high burden, low progressivity
- Treatment 4: high burden, medium progressivity

Randomly exposing subjects to these treatments will provide leverage to discriminate between the hypotheses stated above. By approaching the problem experimentally,

---

<sup>6</sup>The precise distributions of tax rates and benefit shares across pre-tax income groups are displayed below in Table ??

we avoid potential post-treatment biases, and, most importantly, the perennial presence of feedback effects between policy designs and political support for different designs, a problem that limits the scope of what can be gained from approaching the study of compliance on the basis of standard cross-national surveys. Lab experiments are particularly informative because they provide considerable control over ensuring the internal validity of the treatments designed to test our conjectures regarding the factors that shape tax compliance. There are of course concerns about the external validity of experiments that rely on small convenience samples consisting entirely of students. Accordingly we explore the robustness of our lab experiment findings by replicating the experiments online with much larger and more diverse subject pools. These online experiments can provide reassurance that the treatment effects found in the lab generalise to much broader populations.

## 4 Experimental Designs

Our theoretical conjectures are about the micro-foundations of individual reasoning regarding tax compliance. What characteristics of the fiscal regime encourage compliance or, alternatively, generate cheating? The analysis of observational data make it difficult to tease out the precise features of these regimes that shape individual behavior. Accordingly, we have designed both lab and online experiments that help assess our conjectures regarding the causal effects of fiscal regimes on tax compliance.

The broad features of the experiments are the following: each experimental session is conducted employing one of the four different fiscal regimes described earlier; subjects are randomly assigned to groups of four; they earn money from real effort tasks; subjects report the amount they earned and can lie about these earnings if they wish; these earnings are taxed according to the tax rates of the fiscal regimes; there are two different audit rate treatments that determine the probability by which subjects' reported earnings

are verified; deductions applying to the group members are pooled and then shared according to a formula specific to the fiscal regime. Subjects are paid at the end of experiment, and do not receive feedback about earnings until the end of the experiment. One version of these experiments was conducted in the experimental labs of the Centre for Experimental Social Sciences (CESS) at Nuffield College, Oxford. A second version of the experiments was conducted online using the online experimental facilities of the Nuffield College CESS.

## 4.1 The Lab Experiment

The lab experiments were conducted at the CESS experimental lab with 125 University of Oxford students from the CESS subject pool. Participants receive printed instructions at the beginning of each module, and instructions are read and explained aloud.

The tax treatments consist of ten rounds each. Table ?? summarises the treatments that are implemented in these two modules of the experiment. Prior to the tax treatments, participants are randomly assigned to groups of six and we follow a partner matching. Thus, the composition of each group remains unchanged for the two tax treatment modules – in the first tax treatment module, subjects face an audit rate of 0% and in the second the audit rate is set at 10%. Each round of these two tax modules is divided in two stages. In the first module subjects perform a real effort task. This task consist of computing a series of additions in one minute. Their Preliminary Gains depend on how many correct answers they provide, getting 150 ECUs for each correct answer.

We conduct a total of five different sessions that are summarised in Table ?. Note that in each session the tax rate is consistent – what varies across tax modules is the audit rate. Once subjects have received information concerning their Preliminary Gains, participants are asked to declare these gains. A certain percentage or “tax” (that de-

depends on the fiscal regime treatment) of these Declared Gains is then deducted from their Preliminary Gains. In sessions 1 and 3, the group revenues are distributed equally amongst the six participants – these reflect fiscal regimes in which the distribution of social benefits are non-progressive. In sessions 2, 4 and 5, group benefits are distributed in a progressive fashion such that the poorest two group members receive 50 percent of the group revenues; the middle income subjects receive 35 percent; and the richest two are given only 15 percent.

**Treatments: Audit Rates.** There were two audit rate treatments in the experiments: 0% and 10%. In the former treatment the subjects' Declared Gains are not subject to verification; in the latter 10% audit treatment, subjects have a 10 percent probability of being subjected to an audit. If subjects were randomly selected to be audited and if the verification finds a discrepancy between the Preliminary and Declared gains, 50% of the observed discrepancy is deducted from the subject's Preliminary Gains. This deduction applies to the Preliminary Gains and not to the declared amount.

Table 1: Summary of Tax Compliance Experimental Treatments

Session	Subjects	Tax Terciles	AR 1	AR 2	Benefits
1	24	1) 28% 2) 30% 3) 32%	0%	10%	1) Rev/6; 2) Rev/6; 3) Rev/6
2	24	1) 20% 2) 30% 3) 43%	0%	10%	1) 50% Rev; 2) 35% Rev; 3) 15% Rev
3	24	1) 42% 2) 45% 3) 47%	0%	10%	1) Rev/6; 2) Rev/6; 3) Rev/6
4	24	1) 42% 2) 45% 3) 47%	0%	10%	1) 50% Rev; 2) 35% Rev; 3) 15% Rev
5	24	1) 20% 2) 45% 3) 70%	0%	10%	1) 50% Rev; 2) 35% Rev; 3) 15% Rev
Total	120				

In each module there is a certain probability that the Declared Gains are compared with the actual Preliminary Gains in order to verify these two amounts correspond. In one module the probability is 0%, while this probability changes to 100% in the other module. In this essay we only report the results for the zero audit probability treatment. If the audit finds a discrepancy between the Preliminary and Declared gains an extra amount is deducted from the Preliminary Gains. In both modules the amount correspond to 50% of the observed discrepancy. In addition, the regular deduction applies to the Preliminary Gains and not to the declared amount. Deductions applying to the four group members are then pooled and equally distributed amongst those members.

At the end of each round participants are informed of their Preliminary and Declared gains; whether these two amounts have been audited; the amount they receive from the deductions in their group; and the earnings in the round. At the end of each tax module one round is chosen at random, and their earnings are based on their profit for that round. Participants are only informed of their earnings for each tax module at the end of the experiment.

**Heterogeneity** The experiment also measured a set of auxiliary variables designed to allow us to explore heterogeneity in treatment effects. The demographic variables were gender and income. We included a measure of trust and a measure of ideological self-identification. And finally we measured other-regarding preferences with a version of the Dictator Game and measured risk aversion with a standard Holt-Lowry game.

**Dictator Game.** In order to evaluate arguments regarding other-regarding preferences and attitudes about redistributive taxation we included in the first module a Dictator Game. Subjects are asked to allocate an endowment of 1000 ECUs between them and another randomly selected participant in the room. Participants are informed that only half of them will receive the endowment, and the ones who receive the endowment will be

randomly paired with those who don't. However, before the endowments are distributed and the pairing takes place, they may allocate the endowment between themselves and the other person as they wish if they were to receive the endowment.

**Risk Aversion.** Concern about job or status security is hypothesized to shape redistribution preferences. Risk averse subjects should be most enthusiastic about redistributive taxation. The fourth and last module of the experiment consists of a lottery-choice test consisting of ten pairs, which is based in the low-payoff treatment studied in (?). The lottery choices (shown in Table ??) are structured so that the crossover point to the high-risk lottery can be used to infer the degree of risk aversion. Subjects indicate their preferences, choosing Option A or Option B, for each of the ten paired lottery choices, and they know one of these choices would be selected at random ex post and played to determine the earnings for the option selected.

Table 2: Lottery Choices

	Option A	Option B
1	10% of 2.00£, 90% of 1.60£	10% of 3.85£, Bs. 90% of 0.10£
2	20% of 2.00£, 80% of 1.60£	20% of 3.85£, Bs. 80% of 0.10£
3	30% of 2.00£, 70% of 1.60£	30% of 3.85£, Bs. 70% of 0.10£
4	40% of 2.00£, 60% of 1.60£	40% of 3.85£, Bs. 60% of 0.10£
5	50% of 2.00£, 50% of 1.60£	50% of 3.85£, Bs. 50% of 0.10£
6	60% of 2.00£, 40% of 1.60£	60% of 3.85£, Bs. 40% of 0.10£
7	70% of 2.00£, 30% of 1.60£	70% of 3.85£, Bs. 30% of 0.10£
8	80% of 2.00£, 20% of 1.60£	80% of 3.85£, Bs. 20% of 0.10£
9	90% of 2.00£, 10% of 1.60£	90% of 3.85£, Bs. 10% of 0.10£
10	100% of 2.00£, 0% of 1.60£	100% of 3.85£, Bs. 0% of 0.10£

## 4.2 The Online Experiment

In order to explore the broader external validity of the tax compliance experiment conducted in the lab, we implemented a very similar experimental design employing the Amazon Mechanical Turk subject pool. The AMT subjects were informed that they would be randomly matched with a group of five subjects who had already made their decisions. In fact, the AMT subjects were matched to one of the groups that had already played the tax compliance game in the lab under the appropriate fiscal regime characteristics – this was necessary in order to calculate the compensation of each AMT subject.

The treatment assignments in the online experiment closely resembled those of the lab experiment with one major exception: In the online experiment respondents were randomly assigned to either 1) a deadweight loss treatment in which the total group revenues were reduced by a half or 2) a non-dead weight loss treatment in which group revenues were not reduced at all. Hence in the first round of the experiment all subjects (500) were treated with the 10 percent audit treatment and half (250) were randomly assigned to a deadweight loss treatment while the other half (250) had the non-deadweight loss treatment.

## 5 Experimental Results

To test our conjectures regarding tax compliance we first define the dependent variable as the ratio of declared gains to actual earning.<sup>7</sup> A higher ratio implies higher levels of tax compliance. Secondly, we create a variable, "Progressivity", that captures the progressivity of the fiscal regime – both on the tax collection side and in terms of benefits. The variable incorporates information from fiscal regime treatments and assumes

---

<sup>7</sup>There were 26 of the online subjects (less than 3% of the total) who reported income that was higher than their actual earnings – these respondents were excluded from the analysis).



the following values: 4 for the low tax burden and low progressivity treatment; 5 for the high tax burden and low progressivity treatment; 58 for the low tax burden and high progressively treatment; and 40 for the high tax burden and medium progressivity treatment. A third variable in the model measures the level of taxation – it assumes a value of 30 for the low and moderate tax burden treatments and a value of 45 for the high tax burden treatments. In addition in the online models we included control variables for Age, Gender and Left-Right self-placement (Ideology).

Results for the estimated random effects model are presented in Table ???. We included both two-way and three-way interactions designed to assess the extent to which rates of tax compliance vary across the different fiscal regimes for subjects with varying “income” levels. The results for the online experiment are for the treatments in which the audit rates are set to ten percent and the deadweight loss is 50 percent of the group earnings – reasonable representations of actual fiscal regime conditions. The coefficient estimates are difficult to interpret on their own – but they generally support the conjecture that tax compliance falls as levels of progressivity increase and one’s income rises.

Table 3: Percent of Declared Gains Reported

	(1) Online % Declared	(2) Lab % Declared
Progressivity	-0.08** (0.4)	0.01 (0.01)
Preliminary Gains	-0.0001 (0.0001)	0.0001 (0.0001)
Tax Level	-0.01 (0.01)	0.01 (0.01)
Progressivity X Preliminary	0.00 (0.00)	-.0.00001*** (0.000003)
Progressivity X Level	0.003** (0.001)	-0.0002 (0.0002)
Preliminary Gains X Level	0.00 (0.00)	-0.00001** (0.054)
Progressivity X Preliminary X Level	-0.000002*** (0.0000007)	0.0000002*** (0.0000001)
Age	0.002 (0.002)	
Gender	-0.03 (0.04)	
Ideology	-0.02*** (0.01)	
Constant	1.04*** (0.30)	0.53** (0.23)
Random Subject Effects		
Constant	0.31*** (0.02)	0.05** (0.01)
Residual	0.17*** (0.004)	0.36*** (0.01)
N	1250	1558
Subjects	250	72
Wald Chi2	21.64	202.24
Significance	0.02	0.00

Figure ?? and Figure ?? present a first cut of our main results by displaying the marginal effects of pre-tax income and progressivity on tax compliance (measured as the percentage of reported earnings by the survey online participants) for both the lab – here we present the results for both the zero and ten percent audit treatments– and the online experiments. The left panel in both captures how the slope of the level of progressivity changes as income changes. Intuitively, for low income citizens progressivity has a positive effect on compliance. This seems to run counter the bare rationalist approach to the relationship between citizens and the state, the one that suggests rational poor individuals ought to benefit from the state and not comply to maximize their net returns in the short run. Progressivity seems to increase the willingness of the poor to contribute, however marginally, to the commons. Such willingness turns into sharp resistance among higher income citizens: among the latter, progressivity shows a strong negative effect on compliance. Contrary to what one would expect from Rawlsian priors, progressivity leads to lower compliance among middle and, especially, high income citizens. The right panel in Figure ?? and Figure ?? shows the reciprocal marginal effects: how the impact of pre-tax income changes under different progressivity regimes. Under low progressivity, income plays a very marginal role in shaping the level of compliance. As progressivity increases, income becomes a larger, and clearly negative predictor of citizens' willingness to comply.

Figure 3: Simulated Marginal Effects for Lab Experiment Equation

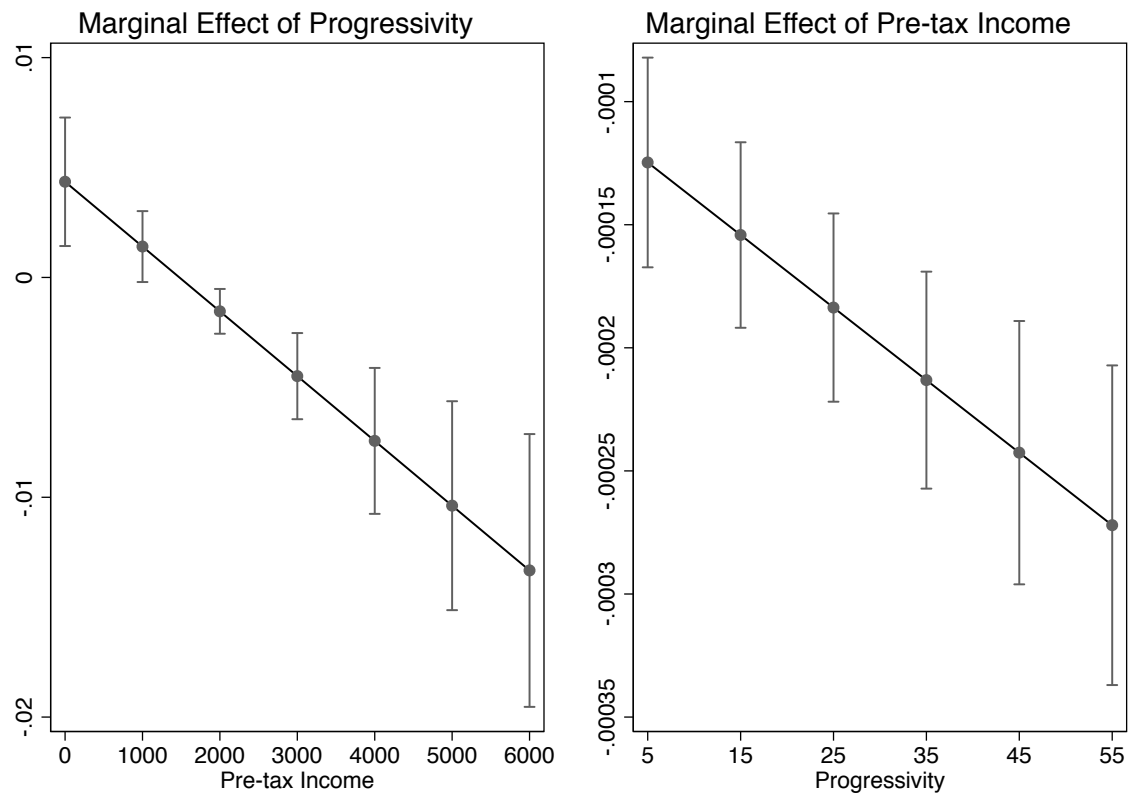
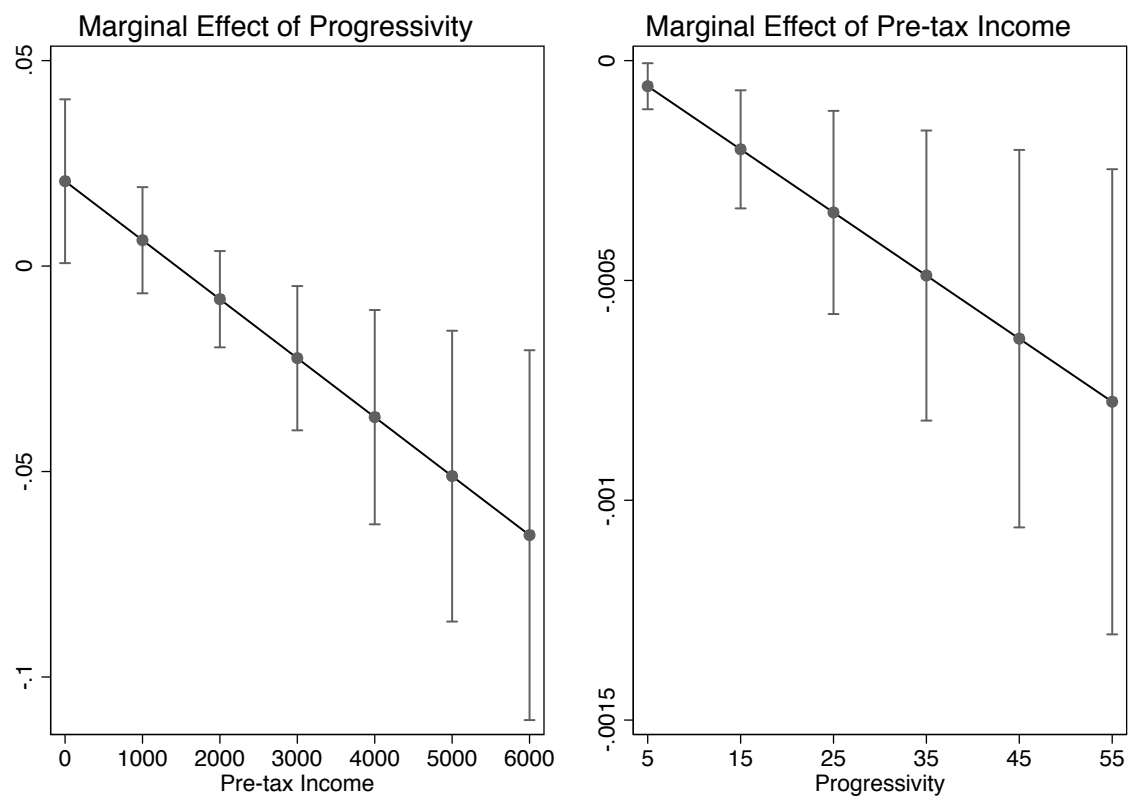


Figure 4: Simulated Marginal Effects for Lab Experiment Equation



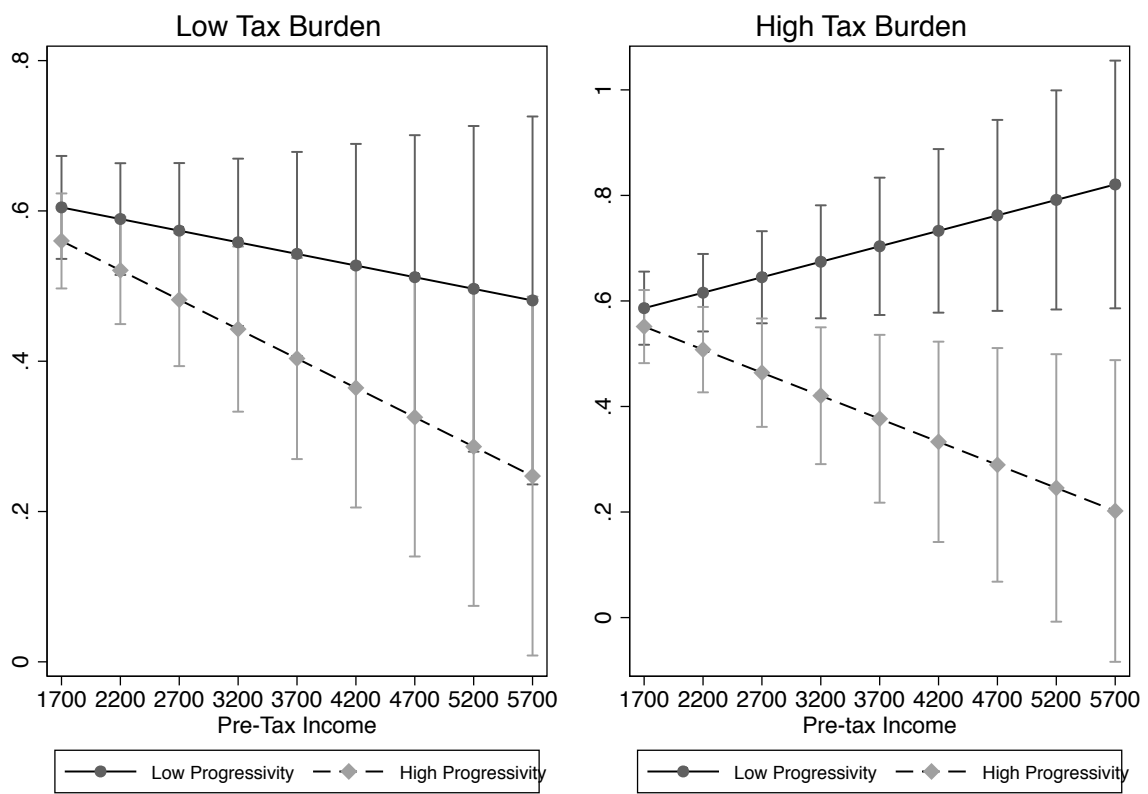
The results in the previous two figures do not condition the role of progressivity or pre-tax income on the size of the fiscal state. Several arguments, referred to above, suggest that links progressivity and the size of the state quite tightly, suggesting that as the size of the fiscal burden increases, progressivity is to be sacrificed (assuming incumbents want to keep a stable political coalition) (??). Moreover, the findings in figures Figure ?? and Figure ??, while clearly providing evidence against the predictions emerging from a Rawlsian framework, could easily be taken as reflecting an income effect. Rich people dislike progressivity, and therefore do not comply. By themselves, Figure ?? and Figure ?? fail to speak directly to the argument about a large common pool from which high income people dip over the long run, thus making compliance self-sustained.

In an effort to explore the viability of this logic more closely, Figure ?? presents the predicted level of tax compliance under different combinations of two regime level variables (the overall level of taxation and the the degree of overall progressivity in the fiscal system, as described above) and the distribution of pre-tax income (that is, the preliminary gains in the online experiment). The results reported present several interesting patterns. First, at low levels of taxation (average tax burden less than 40 %), citizens behave similarly regardless of the progressivity regime. As they do better in the labor market, their pre-tax income increases, and their compliance drops. They do so under high progressivity as much as under low progressivity (note that the confidence intervals overlap in the left panel).

Interestingly, the pattern changes as the overall size of revenues increases (average tax revenues higher than 40%). As the size of the state increases, the contrast between progressivity regimes becomes starker. Under high progressivity, pre-tax income has a strong negative effect on the predicted levels of tax compliance. By contrast, under low progressivity the income and tax compliance appear to be positively correlated. The comparison between low burden low progressivity regimes and high burden low progressivity regimes

suggest that, in the latter, high income people derive a larger share of net benefits from public interventions that translate into higher levels of tax compliance. If the behavior of the rich exclusively reflected an income effect, the observed positive slope for income under low progressivity and high tax burden should not occur. Overall, the pattern of predictions reported in Figure ?? appears more consistent with a “common pool” logic than with a Rawlsian logic of tax compliance. Indeed, the latter has been consistently and systematically falsified by our data analysis. More progressivity consistently leads to less compliance across the board.

Figure 5: Predicted Effects



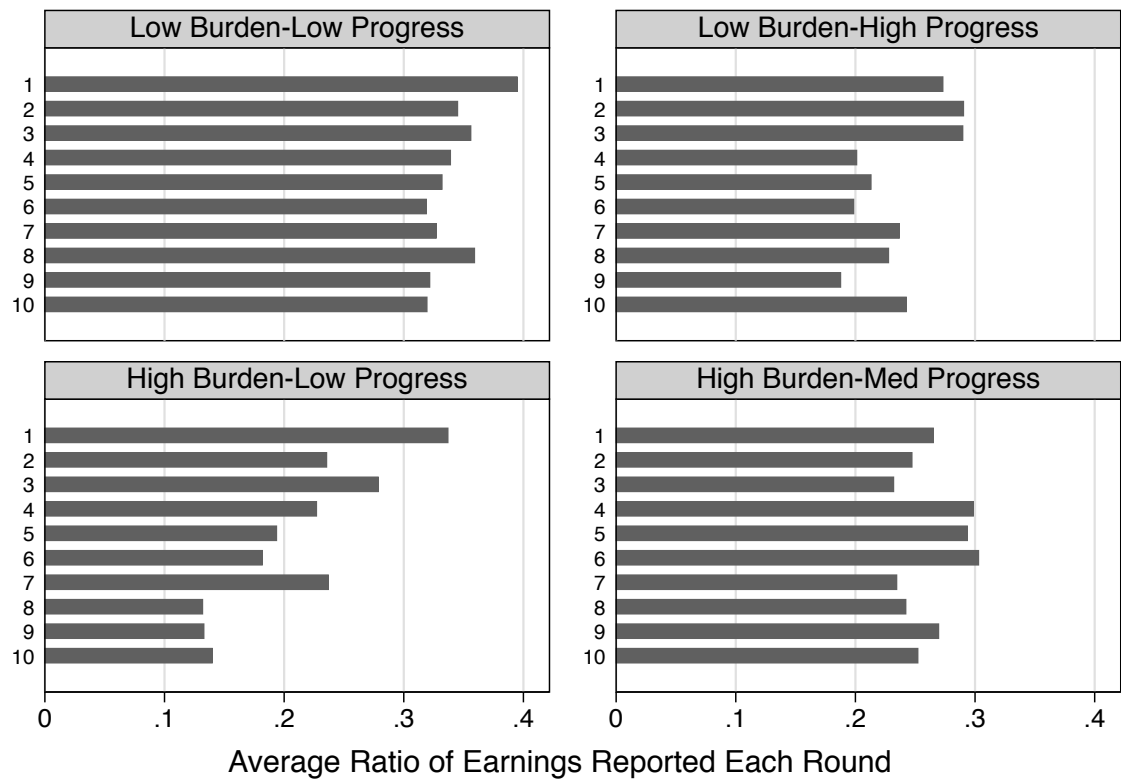


## 6 New Results Section from Ray

### 6.1 Treatments and Results

Figure ?? summarises tax compliance rates across the four tax regimes. Each graph presents the average ratio of earnings reported in each of the ten rounds of the game. Subject behaviour is consistent with our initial conjectures. Focusing on the first row of the graph, note that in the low burden (or low-capacity) condition, compliance is significantly higher in the low progressivity condition (an overall average of 34 percent) compared to the high progressivity condition (an average of 24 percent). The second row of Figure ?? presents the high burden (or high-capacity) condition. Here, again consistent with our initial conjecture, tax compliance behaviour has exactly the opposite pattern to that of the low burden treatment. When taxes are high subjects comply more in the medium progressivity condition (an average of 26 percent) than they do in the low progressivity condition (an average of 21 percent).

Figure 6: Tax Compliance across Regime Types

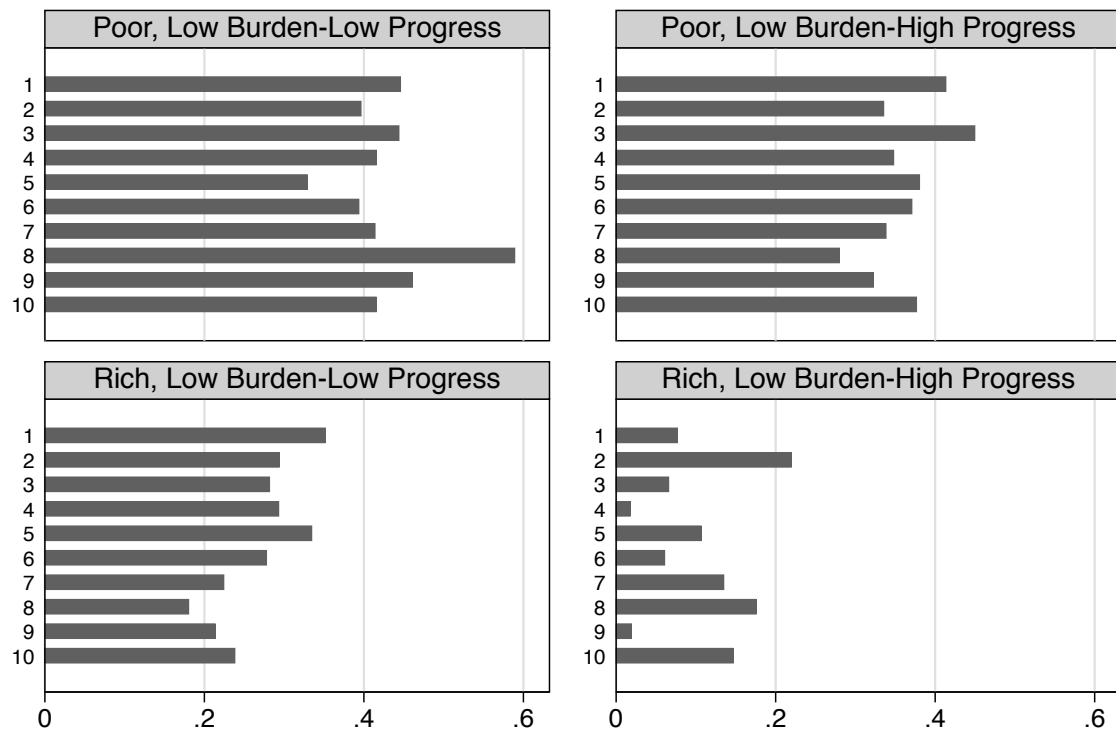


Graphs by regime

Hence when taxes are relatively low, rising levels of progressivity and redistribution result in high levels of tax cheating. In assessing the benefits of redistribution individuals anticipate some dead-weight loss associated with the taxes collected and there is uncertainty as to whether one benefits or loses from redistribution. Weighing this against the relatively small transfers involved, subjects are more likely to retain their earnings by cheating rather than complying with tax policy.

As Figure ?? indicates this effect is strongly driven by the behaviour of the “rich”. Tax compliance by the “poor” responds similarly to high and low progressivity when tax rates are relatively low – their compliance levels are about 35-40 percent. But clearly in a low tax burden context, tax compliance by the rich is very much conditioned on the progressivity of the tax system – under high progressivity the rich report only about 10 percent of their earnings while in the low progressivity treatment they report on average about 25 percent of their earnings.

Figure 7: Tax Compliance for Rich and Poor in Low Tax Treatments



Average Ratio of Earnings Reported Each Round

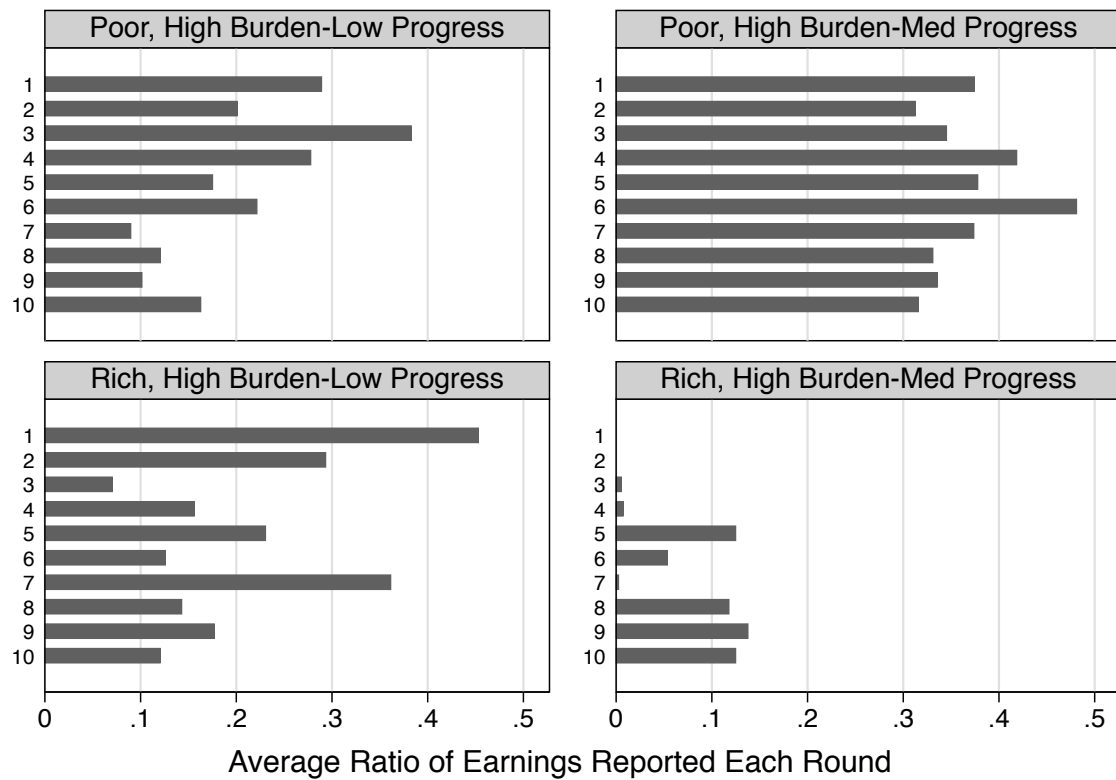
Graphs by income\_type and regime

On the other hand, when taxes are high, rising levels of progressivity and redistribution reduces the rates of tax cheating. In this case the benefits of redistribution to individuals can be quite substantial – and these are payoffs that outweigh any concerns about dead-weight loss and uncertainty regarding one’s type. In this context subjects are more willing to comply with moderate, as opposed to low, levels of progressivity.

This effect, as Figure ?? illustrates, is largely driven by the behaviour of the poor. For the poor in the high tax treatment, higher progressivity results in greater tax compliance. In the low progressivity treatment they report only about 20 percent of their earned income while in the medium progressivity treatment their tax compliance rises to about an average of 35 percent. For the rich, on the other hand, higher progressivity in the high burden tax treatment reduces their tax compliance – it drops from about 15 percent in the low progressivity treatment to about 5 percent in the medium progressivity treatment.

These results suggest that voluntary tax compliance, which is critical for the functioning of the modern state, is very much conditioned on features of the fiscal regime – specifically the overall magnitude of the tax burden which we have characterised as an indicator of state capacity and the progressivity of the fiscal regime.

Figure 8: Tax Compliance for Rich and Poor in High Tax Treatments



Graphs by income\_type and regime

## 6.2 Multivariate Results

Notes on the Results:

**Lab Results** These are the lab results on which I think we should focus. I think they are consistent with the narrative I sketched out above and the coefficients are significant:

low-burden/low-progressivity has higher compliance than low-burden/high-progressivity;  
 high-burden/med-progressivity has higher compliance than high-burden/low-progressivity  
 high-burden has overall less compliance than low-burden the rich consistently cheat  
 strong support for the conjectures

```
.      xtmixed per_declare i.regime ProfitRET level if audit==10 || newssubject:
note: level omitted because of collinearity
```

Performing EM optimization:

Performing gradient-based optimization:

```
Iteration 0:  log likelihood =  -329.0637
Iteration 1:  log likelihood = -328.87471
Iteration 2:  log likelihood = -328.87442
Iteration 3:  log likelihood = -328.87442
```

Computing standard errors:

Mixed-effects ML regression	Number of obs	=	660
Group variable: newssubject	Number of groups	=	60
	Obs per group: min	=	10
	avg	=	11.0
	max	=	20





Performing gradient-based optimization:

Iteration 0: log likelihood = -44.80402  
 Iteration 1: log likelihood = -44.467817  
 Iteration 2: log likelihood = -44.467811  
 Iteration 3: log likelihood = -44.467811

Computing standard errors:

```

Mixed-effects ML regression          Number of obs    =      660
Group variable: newssubject          Number of groups  =       60

                                     Obs per group: min =      10
                                     avg =      11.0
                                     max =      20

                                     Wald chi2(7)      =    174.51
Log likelihood = -44.467811          Prob > chi2       =     0.0000
  
```

```

-----
                per_declare |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-----+-----
                regime |
Low Burden-High Progress |  -.0206446   .0903205    -0.23   0.819    - .1976695    .1563803
High Burden-Low Progress |  -.3845715   .1130539    -3.40   0.001    - .6061532    -.1629899
High Burden-Med Progress |   .6196286   .0991083     6.25   0.000     .4253798     .8138774
                |
                ProfitRET |  -.0001069   .0000268    -3.99   0.000    - .0001594    -.0000543
                |
                regime#c.ProfitRET |
Low Burden-High Progress |  -.0000226   .0000439    -0.51   0.608    - .0001087     .0000636
  
```

High Burden-Low Progress		.0001272	.0000586	2.17	0.030	.0000124	.0002421
High Burden-Med Progress		-.0003234	.0000536	-6.03	0.000	-.0004285	-.0002183
		.3701708	.0575286	6.43	0.000	.2574168	.4829248

Random-effects Parameters		Estimate	Std. Err.	[95% Conf. Interval]	
news subject: Identity					
		.0192158	.0279769	.0011076	.333389
		.2581416	.0073719	.2440899	.2730022

LR test vs. linear regression:  $\chi^2(01) = 0.13$  Prob  $\geq \chi^2 = 0.3599$

**Online Results** I think this is the online result on which we should focus no deadweight loss and audit of 10 the results are weak but consistent with argument in following fashion:

low-burden/low-progressivity has higher compliance than low-burden/high-progressivity;  
 high-burden/med-progressivity has higher compliance than high-burden/low-progressivity  
 high-burden has overall less compliance than low-burden coefficient on earnings is consistent with argument that the rich consistently cheat weak but consistent with lab and theory

```
.
. xtmixed per_declare i.regime prelimGain level if dwl==0 & audit==10 || try:
note: level omitted because of collinearity
```

Performing EM optimization:

Performing gradient-based optimization:

Iteration 0: log likelihood = 35.922754

Iteration 1: log likelihood = 35.922754

Computing standard errors:

```

Mixed-effects ML regression      Number of obs      =      1250
Group variable: try             Number of groups   =       250

                                Obs per group: min =        5
                                avg =        5.0
                                max =        5

                                Wald chi2(4)      =        8.13
Log likelihood = 35.922754      Prob > chi2       =       0.0869

```

```

-----
                per_declare |      Coef.   Std. Err.      z    P>|z|    [95% Conf. Interval]
-----+-----
                regime |
Low Burden-High Progress |  -.0306236   .0558806   -0.55   0.584   - .1401476   .0789004
High Burden-Low Progress |  -.1415646   .0582286   -2.43   0.015   - .2556905  -.0274386
High Burden-Med Progress |  -.0752984   .0593051   -1.27   0.204   - .1915343   .0409374
                |
                prelimGain | -.0000212   .0000181   -1.17   0.241   - .0000565   .0000142
                level |           0 (omitted)
                _cons |   .6726889   .0509397   13.21   0.000   .5728489   .7725289
-----
-----

```

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
-----+-----				
try: Identity				
sd(_cons)	.3128692	.0149069	.2849749	.3434938
-----+-----				
sd(Residual)	.177623	.0039722	.1700059	.1855815
-----				

LR test vs. linear regression:  $\chi^2(01) = 1062.34$  Prob  $>= \chi^2 = 0.0000$

.

. xtmixed per\_declare i.regime##c.prelimGain if dwl==0 & audit==10 || try:

Performing EM optimization:

Performing gradient-based optimization:

Iteration 0: log likelihood = 36.654531

Iteration 1: log likelihood = 36.654531

Computing standard errors:

Mixed-effects ML regression	Number of obs	=	1250
Group variable: try	Number of groups	=	250
	Obs per group: min	=	5
	avg	=	5.0
	max	=	5
	Wald $\chi^2(7)$	=	9.58
Log likelihood = 36.654531	Prob > $\chi^2$	=	0.2139

per_declare	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]	
-----+-----						
regime						
Low Burden-High Progress	-.0979394	.1006295	-0.97	0.330	-.2951695	.0992907
High Burden-Low Progress	-.205389	.1058955	-1.94	0.052	-.4129404	.0021624
High Burden-Med Progress	-.0631129	.1037001	-0.61	0.543	-.2663613	.1401355
prelimGain	-.0000388	.0000361	-1.07	0.283	-.0001096	.000032
regime#c.prelimGain						
Low Burden-High Progress	.0000404	.0000501	0.81	0.419	-.0000577	.0001386
High Burden-Low Progress	.0000375	.0000519	0.72	0.470	-.0000642	.0001392
High Burden-Med Progress	-8.53e-06	.0000514	-0.17	0.868	-.0001092	.0000922
_cons	.7026741	.0736618	9.54	0.000	.5582997	.8470485
-----						

Random-effects Parameters	Estimate	Std. Err.	[95% Conf. Interval]	
-----+-----				
try: Identity				
sd(_cons)	.313396	.0149698	.2853873	.3441536
-----+-----				
sd(Residual)	.1774259	.0039705	.169812	.1853812
-----				

LR test vs. linear regression: chibar2(01) = 1057.10 Prob >= chibar2 = 0.0000

## 7 Discussion

This paper has offered an investigation of how the institutional design of the fiscal system shapes the citizen's levels of compliance. In so far as progressive fiscal designs can be taken to reflect an institutional manifestation of Rawlsian allocative principles, our results strongly suggest that such principles do not guide citizens' decisions about compliance. More progressivity leads to less compliance, particularly in large fiscal states. By contrast, less progressivity in large fiscal states makes high income citizens more compliant with their fiscal obligations. We have posited that this captures a process by which high income citizens see taxes as a long-run investment over their life course.

The results in the analysis, however preliminary, shed some light on a number of important issues in PE of state capacity and redistribution. First, the findings in this paper illustrate a potential mechanism by which high income citizens in large welfare states are contempt despite paying in taxes a large share of their pre-tax income. They also point to the political and institutional conditions that would provoke a potential unraveling of the political coalitions underpinning these large scale redistributive efforts. Second, at the other end of the spectrum the analysis also sheds light on the mechanisms underpinning small states, low compliance traps: in small states, governments are trap between the non-compliance of the rich (if they increase progressivity) or the non-compliance of low-middle income citizens (if they keep regressive fiscal designs that force them to look for functional equivalents in the informal economy). This is particularly the case when low capacity, corruption, and clientelism magnify the deadweight losses incurred by the fiscal state, as both the rich and the poor are pushed towards lower levels of compliance. Indeed, the impact of different levels of *fiscal waste* on compliance across fiscal regimes is an important aspect of the relationship between fiscal institutions and citizens' behavior. We plan to explore this connection more systematically in the future.

## References

- Atkinson, A. B. 1995. *Income and the Welfare State*. Oxford University Press.
- Baldwin, P. 1992. *The Politics of Social Solidarity*. Princeton University Press.
- Beramendi, Pablo and David Rueda. 2007. “The Behavioral Foundations of Social Policy.” *British Journal of Political Science* 37(4):619–641.
- Beramendi, Pablo and Philipp Rehm. 2014. “Who Gains, Who Pays: Progressivity and Preferences.”
- Besley, Timothy and Torsten Persson. 2011. *Pillars of Prosperity*. Princeton University Press.
- Cusack, Thomas and Pablo Beramendi. 2006. “Taxing Work.” *European Journal of Political Research* 45:43–73.
- Dixit, Avinash and John Londegran. 1996. “The Determinants of Success of Special Interests in Redistributive Politics.” *Journal of Politics* 58:1132–1155.
- Doerrenberg, Philipp and Andreas Peichl. 2013. “Progressive Taxation and Tax Morale.” *Public Choice* 155:293–316.
- Duch, Raymond M. and David Rueda. 2013. “The Preferences of the Rich and Poor for Redistributive Taxation: UK and Germany.” Nuffield College CESS Working Paper.
- Fehr, E. and K.M. Schmidt. 1999. “A Theory of Fairness, Competition and Cooperation.” *Quarterly Journal of Economics* 114:817–68.
- Holt, Charles A. and Susan K. Laury. 2002. “Risk Aversion and Incentive Effects.” *American Economic Review* 92:1644–1655.
- Korpi, Walter and Joakim Palme. 1998. “The Paradox of Redistribution and Strategies of Equality: Welfare State Institutions, Inequality, and Poverty in the Western Countries.” *American Sociological Review* 63:661–687.
- Levi, Margaret. 1988. *Of Rule and Revenue*. University of California Press.
- Lindert, Peter. 2003. *Growing Public*. Cambridge: Cambridge University Press.
- Mares, Isabela. 2003. *The Politics of Social Risk*. Cambridge University Press.
- Meltzer, Allan H. and Scott Richards. 1981. “A Rational Theory of the Size of Government.” *Journal of Political Economy* 89:914–27.
- Michelbach, John, John Scott, Richard Matland and Brian Bornstein. 2003. “Doing Rawls Justice: An Experimental Study of Income Distribution Norms.” *American Journal of Political Science* 47(3):530–539.

Rawls, John. 1971. *A Theory of Justice*. Cambridge: Cambridge University Press.

Rawls, John. 1985. "Justice as Fairness: Political not Metaphysical." *Philosophy and Public Affairs* 14:223–251.

Rehm, Philipp. 2015. *Risking Solidarity. Social Policy in Hard Times*. Cambridge University Press.

Varian, H. 1980. "Redistributive Taxation as Social Insurance." *Journal of Public Economics* 14:49–67.