Preferences for Redistribution and Perception of Fairness: An Experimental Study*

Ruben Durante† Louis Putterman‡

November 16, 2009

ABSTRACT

This paper investigates the relative importance of fairness preferences, risk aversion, and self-interest in determining support for redistribution. We present evidence from a series of laboratory experiments in which a large number of subjects choose the level of redistributive taxation to be applied to an initial distribution of endowments among participants. Our design permits us to investigate how support for redistribution varies based on: a) whether or not the decision-maker is part of the group affected by the tax; b) whether or not she has perfect information on her relative position in the distribution; c) whether or not the initial distribution is determined according to task performance; d) on the direct cost of redistribution to the decision-maker; e) on the deadweight loss associated with taxation. We find that: a) most subjects favor a more equal distribution among others; b) support for redistribution is sensitive to the cost of taxation and to the deadweight loss associated with it; c) risk aversion is associated with higher demand for redistribution when income is uncertain; d) subjects support less redistribution when the initial distribution is determined according to task performance. The last effect is much larger for males than for females and accounts for most of the gender-based difference in redistributive choices.

JEL codes: D31, P16, H24, C91

Keywords: income distribution, political economy, redistribution, social preference.

*We are grateful to Roland Benabou, Claudia Biancotti, Samuel Bowles, Jeremy Clark, Pedro Dal Bo, Kfir Eliaz, and Jean-Robert Tyran for very helpful comments. We would also like to thank seminar participants at Brown University, Princeton University, the University of Massachusetts at Amherst, the University of Padua, the University of Trento and the University of Copenhagen, as well as participants at the ESA 2007 Conference and, the 2007 meeting of the Society for the Study of Economic Inequality (ECINEQ) and the 2008 La Pietra-Mondragone Workshop for helpful discussion. We thank Adam Rachlis for his help in initiating the set of experiments that led to this paper, and Gregory Wyckoff for rapid and efficient programming of the software used. Funding for this study was provided by the Alex C. Walker Foundation, the Steven Rattner and P. Maureen White Foundation and the Department of Economics, Brown University.

†Brown University, Department of Economics. Contact: Ruben_Durante@brown.edu

‡Brown University, Department of Economics. Contact: Louis_Putterman@brown.edu.

I. INTRODUCTION

Redistribution of income through government taxes and transfers has long been normal practice in industrial democracies. Using data from the Luxembourg Income Study, Branko Milanovic (2000) estimated that the income share of the bottom two quintiles of households in 14 OECD countries in the early 1990s was on average 14.7% higher when measured on a post-tax-and-transfer than on a pre-tax-and-transfer basis. Even in the U.S., the least redistributive of the wealthy industrialized countries, Milanovic found a difference of almost 8% between the income share of the bottom 40% of households after versus before taxes and transfers.

The question of how much redistribution there ought to be is one that in the end must cross the boundary between positive and normative discussion. But there are many points on which positive economic analysis can be helpful. Studies that attempt to estimate the magnitude of the trade-off between equality and efficiency are one example. An understanding of why income is redistributed can also be pursued as a matter of positive analysis.

Among the possible explanations of why redistribution occurs in democracies is that there is a social consensus behind it, that is, a large majority of citizens feel better off living in a society with less inequality because it reflects their ethical values, increases their perceived personal and property security (Thurow, 1971), or some combination of these or other reasons. In the limit, redistribution could be Pareto improving, i.e. even those with high incomes could prefer some degree of redistribution to occur despite the material cost to them. If redistribution were universally preferred, then an efficient amount of redistribution could in principle be found, whether using the Pareto criterion or by a Benthamite social welfare function.

At the other end of the spectrum of explanations is the possibility that redistribution results from the combination of majority rule and self-interest, as emphasized in traditional political-economic models of redistribution (Meltzer and Richard, 1981; Alesina and Rodrik, 1994; Sinn, 1995, among others). The distribution of incomes in most societies is right-skewed, with the income of the median individual or household being far below the arithmetic mean. Thus, assuming that a given amount of revenue has to be raised by either a head tax (taking a fixed amount per person), a flat tax (taking an equal proportion of income from each person), or a progressive tax (taking a proportion of income that is higher the greater the individual's income), a self-interested median voter will always prefer the flat over the head tax and the progressive over the flat tax, assuming absence of incentive considerations. If government expenditure benefits all more or less equally, political economy models that assume equal participation in elections always predict that progressive taxes will be adopted in market democracies. The same logic can be extended from funding of public goods to provision of health and other services to providing transfer payments. In the absence of incentive and other dynamic considerations, however, such models predict the complete leveling of incomes, something not observed in practice.

The above discussion leaves out one more important reason why self-interest might lead to redistribution: in the absence of adequate means of insuring themselves against negative shocks, individuals with average or above-average incomes may favor redistributive taxation as a form of social insurance (Benabou and Ok, 2001; Alesina and La Ferrara, 2005). For this to happen, tax regimes must be relatively persistent over time, and voters must have some degree of uncertainty about how they will fare in the future. Both these assumptions are reasonable in real contexts.

Our remarks about social or ethical preferences for redistribution are also incomplete insofar as they fail to consider that the value judgments in question may depend on the nature and causes of pre-tax inequality, and on how these are perceived by the voters. Some authors have suggested that differences in voter preferences may depend, at least in part, on their perceptions of whether the distributive outcomes of the market economy

are perceived as fair or not (Piketty 1995; Ravallion and Lokshin, 2000; Graham and Pettinato, 2002; Alesina, et al. 2004; Alesina and Angeletos, 2005; Benabou and Tirole, 2006). Using survey data from several sources, Fong (2001, 2003) finds supporting evidence for the United States that such fairness considerations matter to people.

An extensive experimental literature has investigated how agents' choices in various economic interactions and games of division may be dictated by forces other than self-interest, such as aversion to inequality (Charness and Rabin, 2002; Fehr and Schmidt, 2003; Camerer, 2003) and how the origin of initial entitlements affects the extent of non-self-interested behavior (Hoffman and Spitzer, 1985; Burrows and Loomes, 1994). However, most of these studies have focused on small group interactions, and it is unclear how these findings can be generalized to explain attitudes toward equality and redistribution at the societal level. Our experiment joins a relatively small existing set of studies of preferences on redistribution that are designed with a macro-political economy application in mind and involve choices that are potentially costly in real money terms to the decision-maker (Ackert et al., 2007; Krawczyk, 2007; Beckman et al., 2004; Beck, 1994). Studies in which respondents' statements of preference among distributions have no payoff consequences for them include Amiel and Cowell (1992), Amiel et al. (1999), Johansson-Stenman et al. (2002), and Carlson et al. (2005).

To investigate the extent of social preferences for redistribution, their sensitivity to the determinants of inequality and to perceptions of fairness, and the more general role that self-interest plays in voting for redistributive taxes, we conducted a series of laboratory decision experiments involving a large number of subjects. One goal was to shed light on the degree to which observed redistributive outcomes in democracies are explained by self-interest versus social preferences for equality. Thus, each subject in our experiments was asked to express a preference for redistribution among the micro community of participants both under the condition of being an outside observer of a distribution of income among others, and in the situation of being an affected party with a specific interest stemming from the expectation or knowledge of having a higher or lower pre-tax income. To investigate subjects' willingness to pay for income equalization and their concern over the possibly "leaky" nature of taxation and redistribution, we varied across treatments both a direct cost to the decision-maker and an administrative or efficiency loss to recipients. We also had subjects make decisions both under uncertainty about their relative position in the pre-tax distribution and when uncertainty had been resolved. We used mainly student subjects but also a non-student adult comparison group. Our design contains several new elements, including large group size, explicit replication of an actual country's income distribution, and multiple income determination methods in combination with multiple decision contexts.

We find most subjects willing to pay to increase equality of earnings among others whom they do not know. This willingness varies in predictable ways with the direct cost to the decision-maker, and with their political views and (real world) incomes. It varies in an intuitive way with whether subjects "earn" their unequal laboratory incomes, although this difference itself is sensitive to gender in an interesting manner echoing the political "gender gap" (females are more reluctant to accept even "earned" inequality). Subjects value efficiency, redistributing less when more income is thereby lost. Despite the clear evidence of "social preferences" most subjects' choices regarding redistribution reflect their personal interest when this is also at stake. We also find a link between risk aversion and desire for redistribution. Finally, we show how subjects' decisions can be used to fit utility functions which are clearly concave in the degree of social equality and with which both the median-preferred and the (additive) social welfare maximizing level of redistribution can be calculated. Using these estimates, we compare the redistribution desired by our subjects to the levels of redistribution delivered by real-world democracies.

The remainder of the paper is organized as follows. Section 2 describes the design and rationale of our

experiments. Section 3 provides a theoretical framework for predicting and interpreting the results. In section 4 we illustrate and discuss our main results. Section 5 concludes.

II. EXPERIMENT DESIGN

We designed our experiment to elicit choices with respect to redistribution of income from twenty-one participants in each of sixteen sessions. In outline, each subject first chose her preferred level of a linear tax to redistribute earnings among twenty subjects with pre-tax incomes mirroring the U.S. pre-tax income distribution. Her choice was implemented if he or she became the randomly selected dictator who earned an amount unaffected by it (or largely so) - a disinterested observer condition. Then each subject chose a level of redistributive tax for the alternative situation in which the choosers income was one of the affected twenty. Each choice was in fact a quadruple, with a separate redistributive tax level possible for each of four different determinants of initial incomes, two amenable to interpretation as earned, two not. In about half of the sessions, the second set of choices could be remade after the resolution of uncertainty about own income. Randomizations determined whether the disinterested or interested decision-maker condition held and which of the four methods determined pre-tax income rankings. Although only the eight (or twelve) choices just described are of focal concern to us, subjects required more than 90 minutes to learn the nature of their choices, make the decisions, engage in the tasks potentially determinative of own pre-tax income, make a final set of decisions providing a measure of risk aversion, and complete a background survey. A potential direct cost of redistribution to the decisive individual, and a possible efficiency loss to others, varied among sessions. We now describe the experiment in greater detail.

The sessions began with a set of instructions that appeared on the subjects' computer screens and were simultaneously read aloud by the experimenter so that all subjects were aware of facing identical rules and procedures. At the end of this first instruction stage, subjects were invited to ask questions and then answered five questions to test their comprehension of the procedures. Subjects were informed that there would be two additional parts to the experiment and that further instructions would follow.

As part of the on-screen instructions, we presented a table describing the set of provisional experimental payoffs to be assigned to each of the participants (Appendix Table 1). The distribution of the payoffs, ranging from \$0.11 to \$100.00, reproduced the distribution of the average pre-tax incomes of the lowest to highest earning twentieths of the US population, which was also included in the table.[1] Participants were informed that the provisional earnings might be altered by a tax and transfer process.

In the disinterested observer condition of Part I, each subject was asked to choose a proportional tax rate (0%, 10%, 20%, ..., 100%) to be applied to the pre-tax payoff distribution among the other twenty participants with the proceeds being distributed equally among all subjects. Participants were informed that, at the end of the session, one person would be randomly selected as the "decisive individual," and his preferred tax rate would be applied to the pre-tax earnings distribution of the other twenty participants to determine their final payoff. The decisive individual himself, however, would be affected neither by the pre-tax income profile nor by the tax and transfer to be implemented. By requiring all subjects to indicate their tax preferences at the outset, we aimed at eliciting "outside observer" preferences from the entire subject pool. We used a dictator rather than a median voter design so that subjects would have no reason to vote strategically.

Two additional dimensions of treatment variation were included in order to study agents' willingness to pay

[1] Appendix Table 1's reference to the distribution of income in the United States was partly intended as a framing device, to give decisions a real world macro-economic reference. However, we attempted to steer a middle course, never telling subjects, for example, that "this is an experiment to study your views about the distribution of income," never using words like "just" or "fair," etc. Compare, for example, Frohlich and Oppenheimer (1992) or Johansson-Stenman, Carlsson, and Daruvala (2002).

for a more equal earnings distribution and their concern for aggregate efficiency. The first parameter (which we will refer to as 'tax cost') measures the cost of each additional 10% tax in terms of a direct reduction in the decisive individual's payoff (compare to Andreoni and Miller, 2002). The tax cost parameter could take four alternative values: \$0, \$0.25, \$0.5, or \$1. For example, in a session with tax cost equal to \$0.5, the decisive individual was charged 50 cents for imposing a tax of 10%, \$1 for a tax of 20%, continuing up to \$5 for a tax of 100%. The second parameter measures 'efficiency loss', or the loss in the aggregate payoff of the other participants associated with each additional 10% tax, in line with Okun's (1975) "leaky bucket" argument.[2] This could take three alternative values: 0%, 12.5%, or 25%. For instance, in a session with efficiency loss of 25%, for each \$10 collected as tax, \$2.50 is lost and \$7.50 is divided equally among the twenty affected subjects.

Formally, the post-tax earnings of the twenty affected subjects are given by:

$$\tilde{y}_i = (1-t)y_i + t(1-e)\frac{1}{20}\sum_{j=1}^{20} y_j \tag{1}$$

with y_i being individual i's pre-tax earnings, t being the tax rate chosen by the decisive individual, and e the dead weight loss associated with the tax. The (expected) payoff of the outsider or decisive individual is given by:

$$\tilde{y}_d = y_d - c(10 \cdot t) \tag{2}$$

with y_d being his or her base-payoff, t his preferred tax rate, and c the cost of each 10% tax. Participants were informed that the base payoff of the decisive individual would be randomly drawn from the interval between \$19.80 (the mean pre-tax payoff of the other 20 subjects) and \$21.80.[3] Therefore, the final payoff of the decisive individual was either entirely unaffected by taxes and transfers (when $c =$ \$0, our pure "disinterested observer" benchmark), or else was affected only by the cost of the tax he would choose to impose ("modified disinterested observer" scenario).

Both the tax cost and the efficiency loss parameters were held constant during a given session, allowing their effects to be measured only by between-subject comparisons.[4] The effects of taxation were explained to subjects verbally, graphically, using a table (Table 1), and by means of an equation resembling (1), so that both more and less mathematically inclined subjects could understand them. Subjects were required to pass a comprehension test before making any decision.

Prior to making their Part I choice, participants were also informed that the pre-tax earnings distribution would be determined by one of four possible methods: a) randomly ("Random"); b) based on the average income of their place of origin ("Where From," derived from their home ZIP code, or, for subjects from coun-

[2]The efficiency loss parameter could be interpreted as a measure of the dead weight loss associated with distortionary taxation, or alternatively, as the cost of administering the tax. The latter interpretation was offered to the subjects.

[3]Although it was impossible to totally eliminate comparisons between his own income and that of the other twenty subjects, we chose a base income at least equal to the group average for the decisive individual with the aim of moderating the salience of such concerns. A higher base income would reduce the likelihood of invidious comparisons with higher earners, but increase the likelihood of guilty comparisons with low earners. The impact of the choice of base income can be explored in future experiments. Subjects were told that the identity of the decisive individual would never be revealed, a measure we adopted to eliminate worry over the social tension that he might feel from anyone unhappy with the chosen t. The decisive individual's base income had a random element to make it difficult even for that individual to be sure he had been chosen, again to reduce worries about feelings of tension at the end of the session (this is also the reason why a revision stage was not added in cases in which Part I was randomly selected.). We wanted each subject to focus as much as possible, when choosing tax rates, on the consequences for her and others' earnings, and not on any consequences for their own social interactions with the others at the close of the experiment.

[4]Because each subject already made either eight or twelve tax choices in the session under varied sources of inequality, outsider versus insider conditions, and uncertainty versus certainty of own income, while also performing a number of other tasks, varying tax cost or efficiency within sessions seemed inadvisable.

tries other than the US, from their home country);[5] c) according to their performance on a general SAT-like knowledge quiz ("Quiz"); d) according to their score on a computer-based game of skill ("Tetris").[6] The actual method to be employed would be randomly selected at the end of the experiment. Each subject was asked to choose a tax rate for each of the methods. The four methods were designed to mimic different determinants of economic success in real life (luck, initial conditions, effort and/or ability, respectively) and were used to assess differences in agents' attitude toward redistribution relative to their perception of fairness.

After each subject chose four preferred tax rates for Part I, the nature of Part II was explained, questions were invited, and subjects again took a comprehension test.

Part II was an "involved participant" condition in which each subject was again asked to choose a tax rate for each of the four methods, this time on the understanding that if selected to be the decisive individual, his base payoff would be one of the twenty earnings levels described in Table 1 and his preferred tax rate would be applied to the pre-tax earning distribution among twenty participants, including himself. In this case, another subject was randomly selected to receive \$19.80 to \$20.80 and be unaffected by either the redistribution or the tax cost. This section was designed to analyze the effect of involvement on subjects' choices. The tax cost and efficiency loss parameters varied across sessions but did not vary between Part I and II.

Before choosing Part II tax rates, subjects had to pass another comprehension test. They were then asked to report how they expected to rank under the three non-random earnings determination methods, and how confident they were about their guesses. They then chose the tax rates, took the 20 question Quiz, practiced the Tetris game for two minutes, and played the Tetris game for five minutes. After this, a coin was tossed to determine whether payments would be based on Part I or II. If Part II was selected, participants were informed of their actual ranking in each of the four methods and were offered the possibility of revising their tax choice (we will refer to this stage as 'Part III'). This condition removed subjects' uncertainty about their relative position in the pre-tax distribution allowing us to study the effects of self-interest under certainty and with a wider range of costs than in Part I.[7] Then the earnings-determination method was selected (by the roll of two dice), the decisive individual was chosen (by drawing a code number from a hat),[8] and the final payoffs were announced.

Before exiting the session, subjects were asked to make a series of choices between earning a dollar with certainty and participating in a lottery with a 50% probability of earning nothing and a 50% probability of earning a positive amount which increased from one question to the next (\$1.80 in the first choice, \$2.00 in the second, \$2.33 in the third, \$2.67 in the fourth, and \$3.00 in the last). This is a simple example of the "multiple price list" method of eliciting risk attitudes; see Harrison and Rutstrom, 2008. This section, which was not pre-announced to the subjects, contributed on average an extra \$1.50 (about 6%) to total earnings, and was included in order to generate an indicator of subjects' risk aversion. After completing it, subjects answered a series of background questions regarding their gender, area of study, socioeconomic background, political inclination, and views on inequality and taxation, while cash payments were counted out and brought to them in closed envelopes. The timing of the experimental session is summarized in Figure 1. All the instructions are available at: www.brown.edu/Research/IDE/walkthrough.

Overall, sixteen experimental sessions were held, involving a total of 336 Brown University undergraduate

[5] This information was collected during the log-in procedure, before subjects knew how it would be used

[6] Subjects were told that the version of Tetris to be played was specially modified to put more and less experienced players on a more equal footing.

[7] Whereas the net cost of taxation to the decisive individual ranges from 0 to \$1 in Part I, it ranges from + \$9.3 per 10% tax for the top earner to -\$2.0 per 10% tax for the lowest earner in the revised decision stage. As mentioned in note 3, there was no revision of tax choices if Part I was chosen.

[8] Although subjects themselves had no way to identify code numbers with individuals, this method was used to help convince subjects that the identity of a decisive individual was indeed being determined randomly.

students from a wide range of disciplines. Table 2 summarizes the number of sessions and subjects organized by the exogenous parameters tax cost and efficiency loss. To check the sensitivity of the results to the subject pool, additional sessions were conducted involving a total of 55 adult non-student subjects recruited from the surrounding community. Results of the analysis of these additional sessions are not reported here, but in general they are not qualitatively different than those with students.

III. HYPOTHESES AND PREDICTIONS

In order to predict how subjects will behave in the experiment we need to make some assumptions about their utility functions. A general form for subject i's utility function is:

$$U_i = f(\tilde{y}_1, \tilde{y}_2, \ldots, \tilde{y}_i, \ldots, \tilde{y}_{21}) \tag{3}$$

where $\tilde{y}_{j \neq i}$ represent the post-tax earnings of each of the twenty other participants potentially affected by agent i's decision., and $\tilde{y}_i$ represents i's payoff if he/she is selected as the decisive individual, given by (2).

If individual i is purely self-interested, arguments other than $\tilde{y}_i$ can be ignored without loss of predictive power. Under this assumption, we can predict:

H_{0a}: *In the "disinterested observer" scenario (Part I) a purely self-interested individual will never select* $t > 0$ *if the tax cost is strictly positive (*$c > 0$*). When* $c = 0$ *a purely self-interested agent will be equally likely to select any of the possible tax rates* $(0, 0.1, ...1)$.

In the "veil of ignorance" condition (Part II under random assignment), agent i's choice will depend on the values of c (tax cost), and e (efficiency loss), as well as on his degree of risk aversion. The following hypothesis can be formulated:

H_{1a}: *In Part II under random income determination, a purely self-interested agent will never select* $t > 0$ *if he is risk neutral or risk loving, and if* $c > 0$ *and/or* $e > 0$. *Among risk-averse agents who are purely self-interested, the utility-maximizing* t *is increasing in the degree of risk-aversion and decreasing in* c *and* e.

For the other three methods in Part II, we predict for purely self-interested subjects:

H_{2a}: *In Part II under the Where From, Tetris, and Quiz methods, subjects confident of their predictions about their relative standing will choose 0% or 100% taxation, depending on which maximizes their own expected income. In order to maximize their expected utilities, risk averse subjects lacking confidence in their predictions may select positive tax rates which will be higher the lower the tax cost or efficiency loss, the greater is their degree of risk aversion, the lower is their predicted rank for the method in question, and the lower is their confidence (ability to predict their standing).*

Consider now an individual who, due to social preferences, attaches a positive weight to the earnings of other subjects. We are interested in two types of preferences: preferences regarding equality and preferences regarding efficiency.

Assuming that agents' utility increases with equality in the distribution of incomes (**e**), we can write $h(\mathbf{e})$ as a general function linking utility and equality, with $h' > 0$ if the subject prefers greater equality.[9]

Abstracting from agents' concern for their own income, preference for aggregate efficiency can be formalized in relation to the average of others' aggregate payoffs. Intuitively, the more efficient redistribution is, the larger will be the total pie to be divided among the remaining N_j subjects, *ceteris paribus*. Thus, the utility

[9]The possibility that a subject prefers less equality, especially in cases in which he believes that unequal incomes have been justly earned, will also be considered.

individual i gets from aggregate efficiency can be written as: $g(\frac{1}{N_j}\sum_{j\neq i}\tilde{y}_j)$, with $g\prime \geq 0$.[10]

Formally:

$$U_i = f_i(\tilde{y}_i, x_i) + h_{m,i}(\mathbf{e}, x_i) + g_i\left(\frac{1}{N_j}\sum_{j\neq i}\tilde{y}_j, x_i\right) \tag{4}$$

Function $f_i(\cdot)$ can have varying degrees of concavity, thus incorporating risk aversion, and functions $h_i(\cdot)$, and $g_i(\cdot)$ can vary across individuals both randomly and in relation to a vector of measurable characteristics x_i such as gender, ethnicity, political inclination, and socioeconomic background.

The subscript m in $h_{m,i}$ indicates that i's desire for equality may depend on what method is used to determine pre-tax earnings. For example i may have a strong desire for income equalization under the Where From method if basing earnings on socioeconomic background is perceived by her as unfair, but a much weaker or possibly no desire to redistribute if pre-tax income has been determined by performing a task.

We propose the following compound hypothesis:

H_{0b}. *Both in Part I and Part II (under random income assignment), a subject displaying some level of social preferences may select* $t > 0$ *even if* $c > 0$.

Several sub-hypotheses can be spelled out:

1. The larger c (tax cost) and e (efficiency loss), the smaller the value of t that will be selected, *ceteris paribus*.

2. The greater i's preference for equality under the pre-tax income determination method in question, the larger the value of t the agent will select at every stage, *ceteris paribus*

3. Agents with similar characteristics x will tend to select similar values of t, *ceteris paribus*.

Concern for equality or efficiency does not imply the absence of simultaneously operating self-interest. For example, in both Part II (for any methods other than Random) and Part III, an agent's tax choice will be affected by his expected rank in the pre-tax income distribution via the $f_i(\cdot)$ function. Individuals with higher (lower) expected pre-tax incomes will have a stronger bias toward a low (high) tax. However, concerns for equality and efficiency may have effects countervailing those of self-interest, which will be stronger the closer i's (expected) rank is to the point at which $\partial y_i/\partial t = 0$ (e.g. between ranks 7 and 8, when there is no efficiency loss). Also, since Part II decisions are taken prior to learning one's rank according to the various methods, subjects are expected to prefer higher taxes the greater their lack of confidence in their estimate of their relative performance and the greater their degree of risk aversion. Subjects with (almost) any degree of risk aversion have a self-interested reason to choose a high tax under the Random method, in Part II, if tax cost and efficiency loss are zero (low).

IV. RESULTS

The following analysis is based on the results of the sixteen experimental sessions in which all participants were undergraduate students. Students from a wide range of disciplines participated in the experiment. Subjects were not drawn from particular courses; hence they were not likely to know each other before the sessions.[11] The large majority of participants appeared to have no difficulty understanding the instructions and answering the control questions. Accordingly, all subjects took full part, making tax choices for each of the four methods -

[10] Here, too, $g\prime < 0$ is a possibility, for instance a subject may feel better off the less others earn in comparison to her. We let our data tell us whether subjects value the aggregate earnings of others positively, negatively, or neither.

[11] The 336 subjects were drawn from an undergraduate population numbering about 5700 students at the time of these experiments.

both in Part I and II, and in Part III when this occurred (7 out of 16 sessions). All but one subject also completed the debriefing questions as well as the test for the assessment of risk aversion.

The background questions allowed us to collect information about a number of personal characteristics of the participants. These variables, as well as the risk aversion indicator,[12] are used in the econometric analysis. The distribution of participants by personal characteristics is presented in Appendix Table 2. The questions used to construct the indicators are also reported in the Appendix.

We next illustrate our key findings by presenting the main descriptive statistics. We then discuss the results of a a set of multiple regressions estimated using data from all experimental sessions.[13] The dependent variable - the tax rate selected by each subject - is regressed on a set of explanatory variables which includes: tax cost, efficiency loss,[14] method dummies, risk aversion, a gender dummy variable, ethnic dummy variables, political philosophy, home area income, socioeconomic status, and number of economics courses taken.

Considering the significant share of 0% and 100% tax choices[15], in order to address the concern that, if allowed, some subjects may have chosen a tax rate less than 0% (regressive) or more than 100%, we estimate the regressions using a Tobit model, censored at 0 and 1. We also estimated the same set of regressions using ordinary least squares (OLS) obtaining very similar results. In what follows, we report the results of the Tobit regressions.

*A. **The "disinterested observer" scenario: Part I***

Do agents' tax choices suggest the existence of a demand for redistribution among the micro-community of the twenty other participants? The large majority of subjects display such a demand in the sense that, all things being equal, they prefer earnings to be distributed more equally than the *status quo*, no matter which method is used to determine pre-tax income distribution.Considering all the experimental sessions taken together, in 76.4% of the cases subjects favored some equalization of earnings ($t > 0$), in 44.2% of the cases a tax rate of 50% or higher was chosen, and 14% of the time subjects decided to fully equalize earnings among other participants. The mean tax rate is 42.4%. In principle this result could be due to the choices of those individuals participating in sessions in which redistribution was free or very cheap. However, when only those sessions with a positive tax cost are considered (12 sessions, 251 participants, 1004 tax choices) we observe a very similar pattern. Furthermore, even restricting the analysis to those sessions in which redistribution was more expensive (tax cost = \$1 per 10%) the qualitative result remains the same. A large majority (69.9%) of the participants still opted for a positive level of taxation, more than a third (34.8%) for a tax rate of 50% or higher, and 7.7% were willing to pay a full \$10.00 (approximately half of their expected payoff [16]) to equalize earnings among the other participants. This evidence supports hypothesis H_{0b} against the alternative hypothesis H_{0a}.

Does the existence of widespread support for redistribution imply that agents are not responsive to the cost of taxation? The answer suggested by our Part I data is no. As shown in Figure 2a, participants in sessions characterized by high values of tax cost chose lower levels of taxation than participants in sessions with zero

[12]Of the 335 subjects completing these parts, 308 answered the risk-aversion questions consistently and 27 in an inconsistent fashion, that is they rejected a gamble with high expected value but accepted one with lower expected value. To keep the sample as large as possible, we defined a second measure of risk aversion which could be calculated for both consistent and inconsistent responders. To check robustness, we carried out each piece of analysis also for the restricted sample composed by those who replied consistently. Since the results turn out to be quite similar, we present in what follows, the analysis for the larger sample.

[13]In some cases, we restrict our attention to the sample of tax choices for one of the four methods of pre-tax determination (335 observations). Most of the time, however, we use the larger sample obtained by pooling together all of the 1340 observations (335 subjects by four choices)

[14]Since subjects's choices were very similar for levels of tax cost other than \$1 per 10% (see Figure 2a), in order to simplify the interpretation of the coefficient we use a dummy variable which equals 1 for sessions with tax cost = \$1, and 0 for the others. Similarly for the efficiency loss parameter, we use a dummy which equals 1 for sessions with efficiency loss = 25%, and 0 for the others (see Figure 2b).

[15]A comprehensive description of the distribution of participants' tax choices in Parts I, II and III is reported in Appendix Table 3.

[16]Excluding the show-up fee of \$5.