

The Behavioral Foundations of Social Politics: An Experimental Approach

Benjamin Barber
Pablo Beramendi
Erik Wibbels
Duke University

Abstract

The dominant theoretical approaches at work in the comparative political economy of the welfare state provide two competing accounts for why some governments spend more on social policies than others. In the first, poor voters seek to increase their current income by taxing the rich, and social policy serves to *redistribute* income from the rich to the poor. In the second account, voters seek social insurance against future job loss, and social policy serves as an *insurance* mechanism rather than a redistributive one. Both of these accounts are based on unexamined behavioral assumptions. Most importantly, they share the assumption that voters can clearly distinguish between the redistributive and insurance elements of public policy and, therefore, that individual-level characteristics (income, labor market risks) systematically shape preferences over social policy. Our goal is to examine the soundness of these behavioral assumptions. We do so with an experiment designed to assess preferences for policies that are purely redistributive, purely insurance, and those that combine elements of both. In the experiment, subjects earn money by answering spelling questions correctly and then vote on a tax rate that determines the amount of redistribution and insurance. Our results, though extremely preliminary, suggest that citizens do indeed respond differently to the incentives associated with different combinations of insurance and redistributive components in social policy programs.

We would like to thank Justin Esarey, Timothy Salmon and Charles Barrilleaux for sharing their Ztree code with us.

Why do some governments spend more on social policies than others? The dominant theoretical approaches at work in the comparative political economy of the welfare state provide two competing accounts. In the first, poor voters seek to increase their current income by taxing the rich, and social policy serves to *redistribute* income from the rich to the poor. It follows that political systems in which the poor are more powerful will spend more on social policy (Huber and Stephens 2001). In the second account, voters seek social insurance against future job loss. Because wealthier voters have more to lose from job loss, it is the rich who prefer more spending, and social policy serves as an *insurance* mechanism rather than a redistributive one (Iversen and Soskice 2001; Moene and Wallerstein 2001). It follows that political systems with a higher share of voters exposed to labor market risks will spend more on social policy.

Both of these theoretical accounts are based on unexamined behavioral assumptions. They share the assumptions that voters can clearly distinguish between the redistributive and insurance elements of public policy and that individual-level characteristics (income, labor market risks) rather than cross-national differences (in norms of fairness or labor market institutions, for instance) shape preferences over social policy. The two lines of reasoning differ on whether voters care more about current consumption (the redistribution model) or future consumption (the insurance model).

Our goal with this research is to examine the soundness of these behavioral assumptions. We do so in the context of an experiment designed to assess citizen preferences for redistribution and insurance under various policy treatments that combine different elements of risk, redistribution and insurance. These policy treatments are designed to elucidate whether participants understand the distinction between the redistributive and insurance elements of social policy, and if so how participants evaluate the trade-off between current consumption and future income security. We see such behavioral research as a fundamental precursor to understanding

whether, when, and why each of the contending approaches to comparative welfare states provide valid explanations for cross-national differences in social policies.

As our review of extant work in experimental economics and behavioral psychology demonstrate, there are good reasons to question the behavior underpinnings of both the redistributive and the insurance models of social policy, albeit for different reasons. Most importantly, work in behavioral economics suggests that individuals use rather arbitrary rules of thumb once they are required to think along two dimensions. Since many social policies involve elements of both insurance and redistribution (Moene and Wallerstein 2003), it follows that participants might have a difficult time calculating preferences in a manner consistent with the theoretical predictions in Moene and Wallerstein (2001), Iversen and Soskice (2001) and elsewhere. Or maybe not. The point is that even though there has been considerable behavioral work on redistribution and insurance separately, there has been very little work (actually, none as far as we can tell) evaluating how people evaluate the trade-offs between current and future transfers. Even more important in light of the fact that much social spending combines elements of redistribution and insurance (Moene and Wallerstein 2003), there is no behavioral research on the capacity of voters/citizens to distinguish the distributive implications of policies that have redistributive and insurance characteristics. In short, there is no behavioral research that speaks to the real world of social politics, in which voters must evaluate policies combining elements of both insurance and redistribution.

In the absence of solid behavioral foundations, existing research leaves us with many questions: Are citizens able to distinguish the redistributive and insurance elements of public policy? How do different citizens react to the inter-temporal trade-off between current and future income? How do they perceive and react to social policies that combine different elements of redistribution and insurance? Do low- and high-income citizens share the same perceptions and dispositions across countries? Do high-income citizens differ systematically from the poor in this regard? Drawing on insights derived from our experiment and extensive research in social

psychology and behavioral economics, we begin to answer some of these questions, though we emphasize the importance of additional behavioral research to provide a firmer footing for future research on the political economy of the welfare state.

The paper proceeds in four sections. In the following one, we review the behavioral assumptions underpinning the dominant approaches to the political economy of the welfare state and underscore concerns with those assumptions that emerge from extant research in psychology and behavioral economics. Thereafter we outline our experimental approach to assessing how people evaluate their preferences over redistribution and insurance in a lab setting. The third section offers a set of preliminary findings based on 25 participants who participated in a pilot run. To preview, we find that absolute and relative income are negatively, though weakly, related to preferences for redistribution, that risks seem strongly correlated with preferences for insurance, and that the interplay between income and risk in shaping social policy preferences is contingent on the specific design of the policy instrument—most prominently on the extent to which insurance benefits are targeted exclusively to those who underperform in the labor market. In the concluding section we overview some of the costs and benefits of our experimental approach, suggest avenues for future research and underscore the importance of placing the political economy work on firmer behavioral foundations.

I. The Behavioral Underpinnings of Redistribution and Insurance

Given their prominence, we see little need to provide extended reviews of the redistributive- and insurance-centric accounts of social policy. Power resources theory, the most coherent account of social policy qua redistribution, suggests that where the poor are better organized in encompassing unions and better represented by parties of the left, they have greater success in extracting from the rich and redistributing to themselves (Stephens 1978; Esping-Andersen 1990). In contrast, the varieties of capitalism (VoC) framework emphasizes the importance of specific skills and the demand for insurance against labor market risks as the key

factors underpinning the size of the welfare state. We are less interested in the ongoing debate between these schools of thought than their behavioral assumptions.

The behavioral assumptions underpinning power resources are not terribly explicit. It is unclear, for instance, whether individuals' preferences for redistribution are a function of their perceived position in the income distribution, their own sense of what they merit, their perception of the resources necessary to meet their needs, or concerns for efficiency. Several such motivations could be consistent with the power resources account, but as Scott et al (2001) emphasize, each such consideration is, in fact, quite different. Indeed, they show that all of these considerations enter into voters' preferences over allocative outcomes, albeit with different weights. These different motivations have important implications for the predictions of the theory. For instance, a redistributive model built on citizens' perception of needs rather than their position in the income distribution will often generate very different predictions. While the former would be highly sensitive to the poor's absolute level of income, the latter relies on a voter's position in the income distribution. One model would predict zero redistribution when the poor have a high enough level of income to meet their needs, while the other model would always predict redistribution since every known income distribution is right-skewed.

Certainly the easiest behavioral assumptions for power resource theorists to take on would be those embedded in traditional models of redistribution that emphasize relative income (Meltzer and Richard 1978; Romer 1975). Voters are assumed to be utility maximizers, with utility increasing in income. In a context of two party competition on a redistributive dimension, the median voter will prefer redistribution equal to the difference between their own and the mean income. Overall redistribution should thus reflect the level of income inequality in a society. There are well known problems with this model, the most famous (and hotly debated) being that redistribution seems to increase with equality rather than inequality in a society (Lindert 2004).

All told, empirical confirmation of the Meltzer-Richard model has been hard to come by (Moffitt, Ribar and Wilhelm 1998).¹

Even more important for our purposes, there are reasons to doubt the behavioral foundations of the redistributive model. Indeed, even the simple assumption that voters seek to maximize their own income via redistribution runs into problems. While self-interest clearly plays some role in preferences over redistribution (Esarey, Salmon and Barrileaux 2006; Amiel and Crowell 200X in experiments book),² there is substantial evidence that its role is quite limited. Charness and Rabin (2002), for instance, show that “selfless” behavior can take two forms broadly inconsistent with the Meltzer-Richard model. In some cases, individual concerns for reciprocity and the “poor” are much more important than self interest in shaping preferences over redistribution; in other cases, poor individuals are willing to forego money if they believe that doing so will improve efficiency and overall social welfare. In short, the social consequences of redistribution seem to matter as much, if not more, for experimental participants than the implications of redistribution for themselves. Relatedly, a long line of work on “second dimension” politics and group identity emphasizes the importance of religion, ethnicity and race in conditioning preferences over redistribution (Roemer 1998; Scheve and Stasavage 2006; de la O and Rodden 2008; Austen Smith and Wallerstein 2006), and a cottage industry has developed around the notion that preferences for redistribution are conditioned by an individual’s prospect for upward mobility (Benabou and Ok 2001; Alesina and La Ferrara 2004). In short, we have reason to doubt that voters think and behave as the redistributive framework would seem to imply with regards to redistribution.

¹ See Milanovic (2000) and Finseeras (2008) for exceptions. In the former case, the positive cross-national association between inequality in factor income and cash transfers seems to be driven by pensions. Since pensions are deferred income, this finding would seem inconsistent with a reasonably strict interpretation of Meltzer-Richard. The latter paper uses multi-level modeling to estimate the relationship between individual income, inequality, and individual-level preferences for redistribution in 22 OECD countries and finds evidence supportive of Meltzer-Richard.

² Esarey et al. find that relative earnings are more important than ideology in shaping preferences for redistribution in an experimental setting. Amiel and Crowell find that abstract concerns for the efficiency costs of transfers are strongly attenuated when an individual’s income is at stake.

Equally important, median voter based models of redistribution provide no foundation at all for thinking about the demand for social insurance, the single largest portion of contemporary welfare states. As Moene and Wallerstein (2003: 487) note, more than 30 percent of social policy in the OECD is oriented toward insurance against unemployment, ill health or accidents. Social insurance for the elderly adds very substantially to that share. As Feldstein (2005) notes, the underlying aims and implications of means-tested redistribution and event-induced social insurance are very different. Power resources theory, and the broader model of redistribution on which it rests, provide no theorizing on this substantial share of social budgets, nor the behavioral micro-foundations for thinking about the demand for such policies.

Over the last decade, several prominent papers have forwarded models that integrate labor market risk as a crucial parameter impacting the demand for social insurance (Moene and Wallerstein 2001; Iversen and Soskice 2001). These models integrate considerations of both redistribution and insurance by emphasizing that individual preferences are shaped by their relative income, exposure to risk in the labor market and risk aversion. Because the wealthy have more to lose from job loss or other negative shocks than the poor, there are a wide array of conditions under which the wealthy prefer more insurance than the poor. For Moene and Wallerstein (2001) these conditions entail high levels of risk aversion and universal provision of benefits. For Iversen and Soskice (2001), in turn, risks are a function of asset specificity; thus, high income earners with highly specific assets may prefer higher levels of insurance than average income earners with low levels of asset specificity. The guiding principle underpinning this kind of models is that voters' preferences are not shaped by current income alone but also by the prospect of future income shocks.³

The assumptions upon which these models rest are quite demanding. At the broadest level, we do not know precisely what labor market risk is, how it relates to the related but distinct

³ For the classic formalization on the relationship between insurance and expected utility, see Friedman and Savage (1948), Varian (1980), and Atkinson (1995).

concept of uncertainty, nor how it relates to the full range of insurance and redistributive policies implemented by modern welfare states. Thus, while the insurance framework provides a rigorous theorization of the demand for something like unemployment insurance, it is not clear that it has much bearing on demand for old age insurance, which in most OECD countries consumes much larger shares of the budget. Much more important for our purposes, however, the behavioral foundations of the insurance model are even more suspect than those underpinning the Meltzer-Richard model of redistribution. Even with specific reference to unemployment insurance, the insurance model rests on complex cognition with questionable behavioral foundations. A person's risk assessment involves knowledge not just of their own skills, but of their place in a given firm, that firm's place in a given sector, that sector's position in a broader macroeconomy, and that macroeconomy's place in a broader global economy. The more sophisticated versions of the insurance models also require that each individual conduct the counter-factual analysis of what their job prospects would be in the event that they lost their job. As a large body of evidence makes clear, many voters have a difficult time developing internally coherent positions on a host of "basic" issues (Converse 1964; Zaller 1990), particularly as they bear on issues such as taxing and spending (Bartels 2005).⁴ Indeed, in the context of risk assessment, Shoemaker and Kunreuther's (1979: 616) three-decade old findings that "highlight people's limited abilities to process information, particularly probabilities..." underscore this point, as does the oft-noted tendency of people to over-weight outcomes considered certain (such as, for instance, having a job tomorrow) relative to those that are merely probable (such as the likelihood of having a job six months from now) (Kahneman and Tversky 1979).⁵

⁴ Though see recent evidence in Ansolabehere, Rodden and Snyder (2008) that mass opinions are much more coherent than the conventional wisdom would indicate.

⁵ The experimental insurance literature has also emphasized that once people have to calculate on more than one dimension, as on the trade-offs between premiums and deductibles, they have a hard time making the decisions implied by insurance models. It is not hard to imagine this finding would extend to publicly-financed insurance, on which voters have to calculate on at least two dimensions—their personal tax costs and their expected benefits.

The insurance model runs into other problems from the point of view of behavioral research. For starters, there is some evidence that people view insurance as an investment rather than as protection against future losses (Slovic et al.). Risk aversion, moreover, is not randomly distributed as in Moene and Wallerstein (2001). From groundbreaking work by Kahneman and Tversky (1979) to more recent work by Beckman (2006), it seems that the wealthy are more risk averse and the poor more risk accepting than standard insurance models suggest. More generally, Kleindorfer and Kunreuther (1999) show that risk aversion varies systematically in ways not captured by insurance models. A related body of work has begun to theorize the evidence that responses to risk are highly emotional and fail to respond to the considerations emphasized in the political economy literature on insurance (Loweentson, Weber and Hsee 2001; Edwards, Unigwe, Elaya and Hood 2003). Finally, there is some evidence that labor market vulnerability is as much a function of macroeconomic trends as it is particular skill sets or other individual characteristics as suggested by VoC (Bandyopadhyay and Cowell 2007). If that is the case, it might be that *uncertainty* over the business cycle is as or more important than labor market *risks* emerging from asset specificity in shaping preferences for insurance.⁶ As Starmer (2000) and many others have noted, decision-making under uncertainty and under risk (where the former involves probabilities that are unknown and the latter involves known probabilities) can produce very different outcomes. Together many of these findings make contributions to the growing branches of “non-expected utility theory”, the entire point of which is that the underlying behavioral assumptions of insurance models do not fit the facts.⁷

The underlying accounts of both major theories are complicated by the simple fact that many real-world social policies have elements of insurance and all insurance programs involve varying levels of redistribution. As noted above, several recent models of social insurance

⁶ The distinction between uncertainty and risk is crucial in the behavioral economics literature. The former describes a state in which probabilities are known while the latter describes a state in which at least some of the relevant probabilities are not known.

⁷ See Starmer (2000) for an overview.

explicitly recognize this (Varian 1980; Atkinson 1995; see also Korpi 2006; Iversen and Soskice 2009). The modeling strategies imply that voters can accurately distinguish between the insurance and redistributive elements of social policies and form preferences in a manner consistent with their position in the income distribution, their labor market risks, and degree of risk aversion. But can they? A considerable amount of work in behavioral economics suggests otherwise.

In summary, existing experimental work provides two broad sets of findings that raise important questions about the behavioral foundations of the dominant political economy models of the welfare state. First, it is possible that issues of redistribution and insurance are not independent. Indeed, the study of risk and redistribution are close cousins, both fundamentally concerned with conceptualizing and measuring distributions (Cowell and Cruces 2004). A long line of theoretical work with extensive empirical backing suggests that individual assessments of risk are deeply affected by their position vis-à-vis a reference point, oftentimes the median income. One variant of this thinking is evident in the growing literature on how expectations of upward mobility affect preferences for redistribution (Ravallion and Lokshin 2000; Alesina and La Ferrara 2004; Benabou and Ok 2001). An even larger literature rooted in prospect theory has shown time and again that subjects place excessive weight on very unlikely events and can switch from risk aversion to risk accepting depending on the baseline risk level or their place in the income distribution (Amiel and Cowell 2000; Gonzalez and Wu 1999; Beckman 2006). In either case, evaluations of income and risk distributions are intertwined in a manner inconsistent with dominant models of redistribution and insurance. Second, one common insight in the behavioral economics literature and some branches of political science research is that voters are either not sophisticated enough to appreciate the relationship between taxing and spending or pay insufficient attention to understand the balance between insurance and redistribution which varies so substantially across social policies. The key and very general finding in the behavioral economics and social psychology literatures is that people are easily confused when forced to

think along two dimensions. Forced to think along two dimensions, participants oftentimes rely on rules of thumb that have little bearing on the models of redistribution and insurance.

Yet, while existing experimental work provides grounds for skepticism, it does not deal directly with the issues central to the welfare state literature. Perhaps most importantly, a great deal of experimental work is aimed at very abstract notions of “fairness” as it bears on income distributions and on general assessments of risk aversion or acceptance.⁸ These experiments typically ask participants to assess income distributions from behind the veil of ignorance, their preferences for different payouts under different probabilities, and such. The risk literature is for the most part highly abstracted from real-world assessments of risk in which an individual’s underlying endowments and the broader environment combine to produce shocks with some probability. Oftentimes, the focus is on how individuals’ preferences are shaped by uncertainty, typically defined by placing experimental subjects behind the veil of ignorance. Though informative in a generic way about human psychology, these experiments lack crucial ingredients of the political world we are interested: there is no production, no social choice process, and no redistribution from one person to another. As a result, there is no assessment of the assumptions that underpin social policy models. Those assumptions range from the very crisp distinctions at work in some insurance models to the claim that people, for a whole set of different reasons, are incapable of appreciating and responding to policy subtleties.

In what follows we set out to shed some light on the issue by examining social policy preferences in a laboratory democracy. Our experiment avoids abstract income distributions and abstract risks by tying redistributive and insurance considerations to participants’ skill and effort. We also simulate a political process that redistributes income in a way that allows participants to see the distributive implications of transfers in a clear way. Our experiment also goes beyond

⁸ There is also a narrow experimental literature aimed at assessing when and why people will buy private insurance.

existing experimental work by combining concerns with risk and redistribution.⁹ It goes without saying that a substantial body of work indicates that concerns over fairness, social welfare, and the like are more important for preferences over social policy than an individual's risk exposure or position in the income distribution. That said, very little of the work in behavioral economics or social psychology examines the key distinction between insurance and redistribution that lies at the heart of the contemporary welfare state literature. Given the tendency for social policies to combine elements of insurance and redistribution and the theoretical models suggesting that individuals can calculate their preferences along both of these dimensions, our experiment should speak more directly to the politics of the welfare state than others we find in the literature.

II. The Experiment

This experiment aims to resemble society with a production (market) and a redistribution (policy) phase. We expose participants to three policy configurations: a redistribution wave, an insurance wave, and a hybrid of the two. In turn, each phase consists of three parts: a voting phase in which the median voter decides a tax/insurance rate, a production phase in which participants earn money by answering spelling questions correctly, and a phase in which they view the outcome of their effort and the tax-and-transfer process. In the production phase, participants are provided multiple-choice options for spelling 20 words. Correct answers earn \$.16. Wrong answers are penalized \$-.04, thereby ensuring that guessing is not rewarded. The production phase draws on Salmon et al (2006: 3), who emphasize that the spelling task “mirrors aspects of the real-world work environment that we want to parallel: the task is difficult, uninteresting, performed under deadline pressure and some people are intrinsically more skilled

⁹ Even in experiments in which researchers purport to be examining both considerations, such as in Amiel and Cowell (2000) and Cowell and Cruces (2004), the insurance and redistribution treatments are entirely separate. This separation is justified as necessary to prevent bias, whereby a similar set of questions on risk and inequality to the same respondent might bias their answers. This problem seems particularly relevant to experiments that are based on evaluations of abstract income or risk distributions. As described below, our experimental approach attempts to resolve this problem by making income and risk distributions concrete.

at it than others.” Even in our preliminary test sample, which is significantly weighted toward individuals with high levels of education, we found considerable variation in participants’ performance.

We describe each phase to the participants in detail and with examples to ensure that they understand the workings of the experiment. In order to familiarize themselves with the basic production process, participants have one trial period to answer spelling questions without earning any money. They are then shown how many practice questions they got correct, how much they would have earned, and how all the other participants did. Thereafter, we introduce the mechanics of the political process to participants, where the political process involves voting on a tax rate. The tax rate is determined by the median voter. To familiarize themselves with the political process, participants have one practice round in which each participant votes, does the spelling exercise, and sees the tax rate chosen by the median voter, their earnings before taxation, their earnings after taxation, as well as those of all the other research subjects participating in the experiment at the same time.

After the practice round, the participants move to the first phase, in which government policy is purely redistributive. This phase consists of three parts: First, the participants vote on their preferred tax rate; second, the production phase takes place; and third, the participants see how many questions they answered correctly, how much they earned, and how much each of the other participants earned. Participants also see the tax rate as determined by the median voter, and how much they have earned after taxation. This process repeats itself three times.

Following the regular redistribution phase, the participants are introduced into an insurance phase that is meant to reflect the dynamics of unemployment. In this phase, the participants are randomly assigned a risk of getting very hard questions. These questions are so hard that receiving them ensures very low income. First the participants are introduced to the concept of risk and insurance. The risk of ending up in the bad state of the world (very hard words) ranges between 0 and 0.5. This choice is informed by empirical estimates of the actual odds that a

worker becomes unemployed in the United States.¹⁰ We explain that the tax rate in this case will be applied to each participant but that all taxes will go into an insurance pool that is paid out only to those who get hard questions. Once again, the insurance rate is determined by the median voter. Each participant is shown their randomly assigned risk of getting hard questions in the next phase and asked to vote on the insurance rate. As participants move into the production phase, they receive hard questions conditional on their initial probabilities—higher risks are more likely to receive hard questions. Then the participants answer questions to earn money, unaware of the type of questions they are receiving. At this point, participants see how many questions they got correct and their gross earnings. They also see the insurance rate as decided by the median voter, the implications of the insurance tax on their gross income, and the pre- and post-tax income of all the other participants. Lastly, the participants see how many people received the hard questions and how much those people earned before and after the insurance payout. This process is repeated three times.

The final wave combines elements of redistribution and insurance. The participants are told that half of the tax money will be paid equally to every individual (redistribution), and the other half will be paid exclusively to those participants who got the hard questions (insurance). This 50-50 split is based on a recent stream of research showing that the distributive implications of social spending across most countries of the OECD are neutral (Immerwoll et al., 2005). As in the pure insurance portion of the experiment, participants are shown their risk of getting the hard set of questions and asked to vote on their preferred tax rate. Then, the production phase starts, with participants answering questions to earn money. When all participants finish answering questions, they learn how many answers they had correct and how much money they earned from that. They also learn the tax rate chosen by the median voter, how much money they earned after taxes, how much they earned (or lost) from the redistributive aspect of taxation, and how much they earned (or lost) from the insurance aspect of taxation. As in the pure insurance phase,

¹⁰ We thank Philipp Rehm for his input on this point.

participants are informed as to who had the hard questions and how much they earned before and after taxes. This process is repeated three times so that participants can familiarize themselves with the process.

At the end of the experiment, each participant receives payment of \$15 for participating, plus any additional money earned during the non-practice phases of the experiment. This additional payment reflects each participant's performance in the production phase and the extent to which the individual benefited from the policies simulated in the treatments. The maximum earnings in our initial runs was \$34.21, and the minimum was \$23.19.

Existing research provides clear expectations as to what we should find in each of the three phases of the experiment. In the redistribution phase, poor spelling performers (the poor) are expected to want higher taxes than good spellers (the rich). In the insurance phase, we expect those with the highest risk of very hard questions (the bad state of the world) to prefer the highest tax rates for the insurance pool. Poor participants (bad spellers) should favor less insurance-oriented policies while rich participants should prefer the opposite in light of the decreasing marginal returns of current consumption, though we are unsure whether or not the sums involved in our experiment will evoke such behavior. The most interesting treatment will be the one in which participants are primed on social policies that combine elements of redistribution and insurance. The existing political economy literature suggests that participants will be able to distinguish the redistributive and insurance elements of the tax and form a preference over the policy reflective of their income and risks.

III. A Preliminary Empirical Illustration

We ran a test run of the experiment on 25 participants, many of them associated with the Duke Political Science Department. The results we present here are only meant to serve as an example of the kind of analysis we will be able to conduct when the experiment is run in earnest in June of 2009.

Turning first to the pure redistribution portion of our experiment, Figure 1a plots the tax rates each participant voted for against their net earnings in the previous round, Figure 1b plots the voted for tax rate against the difference between each participants net earnings and the mean for that round. The patterns are not exactly crystal clear, but they do show a negative relationship between both absolute earnings (Figure 1a), relative earnings (Figure 1b), and preferred tax rates. The bivariate correlation is $-.34$ between absolute earnings and preferred tax rates and $-.28$ between relative earnings and preferred tax rates. As Figure 2 shows, the preferred tax rates among those above and below average earnings in both usable rounds of the redistribution game are quite different. Those earning above the mean voted for, on average, a tax rate of 31.7 and 27.5 percent in rounds 3 and 4, respectively. Those earning below the mean voted for tax rates of 40 and 43 percent in each of those rounds, on average.

[Figure 1 and Figure 2 Here]

These correlations, though, do not approach standard levels of statistical significance. This is hardly surprising since we are dealing essentially with 20-25 individuals over three rounds of the experiment. However, even under these rather extreme conditions, some interesting, and again extremely preliminary, findings emerge from the other two rounds of the experiment. The exploitation of the data emerging from all three waves is reported in Table 1. There, we analyze by wave the determinants of the tax rate chosen by the participants in the different rounds. The IV of interest is the level of pre-tax earnings obtained in the previous round, and, in rounds II and III, the risks faced by each individual before choosing the tax rate. Recall that this risk is defined by the probability that each individual ends up in the bad state of the world, that is to say the odds that he/she receives an extremely hard set of spelling words. All models in Table 1 report OLS estimates with robust standard errors clustered by cross-sectional unit (i.e. *personid* in this case).

[Table 1 here]

Before turning to the results, two features of the analyses in Table 1 deserve commentary. First, we have added two obvious controls. These are measures of the ideological profile and risk

aversion of each participant.¹¹¹² We have also added a period dummy (round) to account for any unobserved effect due to the fact that the policy treatments are experienced sequentially. Since all participants have similarly high levels of education, there is no explicit control for that in the regressions. Second, two different model specifications are reported for waves II and III. Model 1 assumes that the effects of earnings and risk are independent from one another. In contrast, model 2, in line with the assumptions in the theoretical literature (Moene and Wallerstein 2001; Iversen and Soskice 2001), is estimated on the premise that the effects of income and risk are conditional on one another. We turn now to discuss the main findings.

Note first that the hard questions did their job; those who fell in the bad state of the world earned, on average, $\frac{1}{4}$ of those who did not receive the treatment. Thus, it is hardly surprising that the measure of risk emerges as a very strong, positive, and statistically significant predictors of the preferred tax rate. Consistent with previous findings, higher exposure to risk increases the demand for insurance (see Rehm 2008 for a detailed discussion and similar results using alternative methodologies). Somewhat less obvious results emerge from the contrast between waves II and III of the experiment. Recall that the key difference between them is that in wave II the pool of insurance benefits is distributed only among those respondents that did in fact end up in the bad state of the world. Social policy is purely an insurance device from which people only benefit after suffering a decline in their earnings due to a bad turn of events in their lives. Wave III, in contrast, comes closer to a universalistic social policy system in which social benefits are distributed among all contributors, combining an insurance and a redistributive dimension.

¹¹ The ideological profile of the respondent is captured through a four point scale that captures the reaction of the respondent to the following statement: “On the whole, do you think it should or should not be the government’s responsibility to reduce income differences between the rich and the poor”. The values of the scale are as follows: 0-“definitely not” 1-“probably should not” 2-“probably should” 3-“definitely should”.

¹² The index of risk aversion is simply the average of the five items included in the standard questions on risk aversion performed by the Panel Study of Income Dynamics. We thank Philipp Rehm for his suggestions on this point. Source and details available from (<http://psidonline.isr.umich.edu/Data/Documentation/Cbks/Supp/rt.html>) Positive answers (“Yes”) are recoded as 1. Negative answers are recoded as 0. Values closer to 0 reflect very high levels of risk aversion. Values closer to 1 characterize risk-taking respondents.

Consider Model 1 first. The negative effect of earnings on the preferred tax rate is significantly different from zero only when the insurance benefits are targeted to the people that performed poorly in the market. While one should not make much ado about issues of statistical significance when working with forty observations, the contrast between waves II and III in this regard is consistent with the idea of a tax base effect: in wave II, high earners are net contributors and know they will get nothing back through the welfare systems. In wave III, high earners are also contributors, but they expect to get back some of their contributions to the fiscal system. This introduces a second motivation that makes their position less clear-cut, and could be interpreted as suggesting that when people who pay in also get their per capita share, they might not necessarily be opposed to taxes.

The contrast between waves II and III of the experiment becomes even sharper when one compares the non-linear specifications (Model 2). The upshot of the results is that the way in which benefits are distributed shapes the way respondents evaluate risks and establish their preferences over social policy. As the coefficients in Table 1 make little sense at face value, Table 2 presents a comparison of the predicted choice of tax rates by our respondents in waves II and III, displaying the range of possible combinations of risk and pre-tax earnings in the previous round.

[Table 2 here]

A number of interesting, though again extremely preliminary, insights emerge from this exercise. Most interesting among them is the balance between the income and the risk effects under different institutional configurations. In a world with insurance benefits targeted to poor market performers, *the income effect clearly dominates the risk effect*. This is reflected by the fact that in the upper half of table 2, at all levels of risk incidence, higher levels of earnings in the previous round lead the subjects of our exercise to choose lower tax rates. Respondents with a low chance of entering the bad state of the world (0.1) go from a tax rate of 45 (lowest level of earnings) to a tax rate close to zero (highest level of earnings). Likewise, while a subject with a .5

chance to enter the bad state of world will choose a tax rate of 100 percent, her choice would drop to 27.4 if she became the best performer in the market. In other words, under the specific institutional conditions specified in wave II, insurance is never a normal good, regardless of citizen's levels of risk aversion.

The results for wave III (lower half of table 2) offer a very different picture. Whereas for lower levels of risk incidence, the income effect continues to prevail, *for sufficiently high levels of risks incidence (>0.2), insurance does become a normal good*, that is to say, the higher the pre-tax earnings of respondents in the previous round, the higher the preferred tax rate. While respondents with a fixed level of risk of 0.2 show a negative elasticity with respect to income in their choice of tax rates (the predicted choice ranges from 48.27 for person with the lowest level of earnings to 17.2 for the person with the highest), the pattern reverses for people exposed to higher levels of risk. For instance, given a risk level of 0.35, a poor performer would choose a tax rate of 45.30, whereas a high wage earner would choose a tax rate of 64%. The fact that benefits are no longer targeted to the worst performers of the market seems to have a strong impact in the way our pool thinks about social policy.

To summarize, we find that, at least for this pool of participants, the way people think about tax policy responds systematically to different combinations of three factors: the level or pre-tax earnings, the exposure to risk, and the extent to which the benefits pool are distributed according to either (1) targeted non-redistributive insurance criteria or (2) comprehensive criteria combining elements of insurance and redistribution. Only under the latter institutional conditions does the notion that insurance is a normal good (Moene and Wallerstein 2001) receive support from our preliminary results. Our findings also suggest that the subjects engaging in this pilot run appear responsive to different structures of incentives. The extreme skepticism about people's ability to differentiate, evaluate, and react to the details of different policy programs appear unwarranted in light of this admittedly crude first cut. Whether these findings will actually hold

up when the experiment is run on a properly randomized selection of subjects remains an open issue.

IV.-Conclusion

Clearly there is much work to be done. The truly comparative goal is to understand how citizens evaluate the trade-off between current consumption and future security across space, institutional contexts, and the life-course. Only thereby will we be able to map the landscape of human dispositions onto different types of social policies. To do so will obviously require a much larger, cross-national experimental and survey-based effort. Doing so will place existing research on sounder behavioral foundations and contribute to the ongoing attempt to mediate between the contending theoretical explanations for the size of the welfare state. Short of a massive cross-national experimental project, we can imagine a number of modifications of our experiment that would allow us to tap into different aspects of redistribution and insurance. For instance, while our experiment introduces riskiness in future income, it does not get directly at the concept of asset specificity. The latter, as developed by Iversen and Soskice, suggests that a negative income shock in time t is likely to be associated with long-term unemployment and reduced income at $t+1, t+2...$ in the absence of retraining policies. We believe we can capture such dynamics by assigning people risks at the beginning of a phase and requiring that they live with it across the phase rather than having a new level of risk reassigned with each production phase. We also might manipulate the fairness of pay. In this experiment, all participants are paid equally for correct answers, but as Salmon et al (2006) have done, we could alter the returns to correct answers so that better spellers “earn” more to see how different pay distributions condition preferences for insurance and redistribution. It would also be quite plausible to manipulate the design in a way that would speak to the growing economics literature on the relationship between upward mobility and redistribution.

It is possible that individuals need not be capable of calculating the complex shades of redistribution and insurance in the real world of social policy for the overarching outlines of policy to be consistent with the insurance models outlined above. Most plausibly, it might be the case that political parties solve these information problems for voters in a manner consistent with Aldrich (1995), Ferejohn (1999) and others. In this account, individuals need not have clear ideas about their risk profile, their place in the income distribution or the relevant policy tools. They simply need a vague notion that one party or the other prefers more or less redistribution or insurance spending. Our experiment does not speak at all to this possibility.

Nevertheless, such possibilities underscore the need for additional work to pin down the how and why of individual-level preferences for insurance and redistribution. Standard caveats of experimental research apply to everything we have done here. Most importantly, we believe experimental research has exactly zero external validity. As such, our experimental results will only be as convincing as the ancillary data that will come from related data collection efforts. We have placed a battery of questions relevant to this project on the most recent CCES survey, the results of which await analysis. We also have a proposal before the CSES to include a battery of questions in the next wave that will help us further parse how citizens understand inequality and risks. We are perfectly aware that our experimental results would be greatly strengthened, if they were complemented by other types of data. Against these considerations, however, we must weigh the well-known difficulty of pinning down causal inferences in the context of observational data—a difficulty particularly salient for a literature that increasingly relies on complex, interactive and multi-level models with multiple feedback loops in the context of 15 or so cases.

References

- Alesina, Alberto and Eliana La Ferrara. 2004. "Preferences for Redistribution in the land of opportunities." *Journal of Public Economics* 89: 897-931.
- Aldrich, John. 1995. *Why Parties?: The Origin and Transformation of Political Parties in America*. Chicago: University of Chicago Press.
- Amiel, Yoram and Frank Cowell. 2000. "Attitudes Toward Risk and Inequality: A Questionnaire-Experimental Approach. DARP (Distributional Analysis Research Programme) Discussion Paper #56. London School of Economics.
- Ansolabehere, Stephen, Jonathan Rodden, and James Snyder. 2008. "The Strength of Issues: Using Multiple Measures to Gauge Preference Stability, Ideological Constraint, and Issue Voting." *American Political Science Review* 102: 215-232.
- Atkinson, Anthony B. 1995. *Incomes and the Welfare State*, Cambridge University Press.
- Austen-Smith, David and Michael Wallerstein. 2006. Redistribution and affirmative action. *Journal of Public Economics*. 90(10-11): 1789-1823
- Bandyopadhyay and Cowell. 2007
- Bartels, Larry. 2005. "Homer Gets a Tax Cut: Inequality and Public Policy in the American Mind." *Perspectives on Politics* 3: 15-31.
- Beckman, Steven. 2006. "A tax and redistribution experiment with subjects that switch from risk aversion to risk preference". *Social Choice and Welfare* 27: 627-641.
- Benabou, Roland and Efe Ok. 2001. "Social Mobility and the Demand for Redistribution: The POUM Hypothesis." *Quarterly Journal of Economics* 116: 447-87.
- Charness, G. and Matthew Rabin. 2002. "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics* 117: 817-869.
- Converse, Philip. 1964. "The Nature of Belief Systems in Mass Publics." In *Ideology and Discontent*, ed. David Apter. New York: Free Press.
- Cowell, Frank and Erik Shokkaert. 2001. "Risk Perceptions and Distributional Judgements." *European Economic Review* 45: 941-52.
- Cowell, Frank and Cruces. 2004.
- Cusack, Thomas, Torben Iversen and Philipp Rehm. 2006. "Risks at Work: The Demand and Supply Sides of Government Redistribution." *Oxford Review of Economic Policy* 22: 365-89.
- De la O, Ana and Jonathan Rodden. 2008. "Does Religion Distract the Poor?: Income and Issue Voting Around the World." *Comparative Political Studies* 41: 437-76.

Edwards, Unigwe, Elaya and Hood 2003

Esarey, James, Timothy C. Salmon, Charles Barrilleaux. 2006. "What motivates Political Preferences? Self-Interest, Ideology, and Fairness in a Laboratory Democracy." Working Paper, Florida State University Department of Economics.

Feldstein, Martin. 2005. "Rethinking Social Insurance." *American Economic Review* 95:1-24.

Ferejohn, John. 1999. "Accountability and Authority." In Adam Przeworski, Susan Stokes, and Bernard Manin, *Democracy, Accountability, and Representation*. New York: Cambridge University Press: 55-97.

Finseeras, Henning. 2009. "Income Inequality and Demand for Redistribution: A Multilevel Analysis of European Public Opinion." *Scandinavian Political Studies* 32: 94-119.

Fong 2001.

Friedman, M and L. Savage. 1948. "The Utility Analysis of Choices Involving Risk." *Journal of Political Economy* 56:

Fong, Christina. 2001. "Social Preferences, Self-Interest and the demand for redistribution." *Journal of Public Economics* 82: 225-246.

Gonzalez, Richard and George Wu. 1999. "On the Shape of the Probability Weighting Function." *Cognitive Psychology* 38: 129-66.

Immervoll, H. 2005. "Household Incomes and Redistribution in the EU: Quantifying the Equalizing Properties of Taxes and Benefits" IZA-Working Paper.

Iversen, Torben and David Soskice. 2001. "An Asset Theory of Social Policy Preferences." *American Political Science Review* 95: 875-893.

Iversen, Torben and David Soskice 2009. "Distribution and Redistribution: The Shadow of the XIX Century" *World Politics*, forthcoming.

Kahneman and Tversky 1979

Kleindorfer, Paul and Howard Kunreuther. 1999. "The Complementary Roles of Mitigation and Insurance in Managing Catastrophic Risks." *Risk Analysis* 19: 727-738.

Korpi, Walter. 2006. "Power Resources and Employer-Centered Approaches in Explanations of Welfare States and Varieties of Capitalism: Protagonists, Consenters, and Antagonists" *World Politics*, 58/2, p. 167-206

Loewenstein, George, Elke Webe and Christopher Hsee. 2001. "Risk as Feelings." *Psychological Bulletin*: 267-286.

Meltzer, Allan and Scott Richard, 1981. "A Rational Theory of the Size of Government." *Journal of Political Economy* 89: 914-27.

Michelbach, P., J. Scott, R. Matland, and B. Bornstein. 2003. "Doing Rawls Justice: An Experimental Study of Income Distribution Norms." *American Journal of Political Science* 47: 523-539.

Milanovic, Branko. 2000. "The Median-voter Hypothesis, Income Inequality, and Income Redistribution: An Empirical Test with the Required Data." *European Journal of Political Economy* 16: 367-410.

Moene, Karl Ove and Michael Wallerstein. 2001. "Inequality, Social Insurance, and Redistribution." *American Political Science Review* 95: 8..-874.

Moene, Karl Ove and Michael Wallerstein. 2003. "Earnings Inequality and Welfare Spending." *World Politics* 55: 485-516.

Moffitt, Robert, David Ribar, and Mark Wilhelm. 1998. "The Decline of Welfare Benefits in the U.S.: The Role of Wage Inequality." *Journal of Public Economics* 68: 421-52.

Roemer, John. 1998. "Why the Poor do not Expropriate the Rich: An Old Argument in New Garb." *Journal of Public Economics* 70: 399-424.

Romer, Thomas. 1975. Individual Welfare, Majority Voting and the Properties of a Linear Income Tax." *Journal of Public Economics* 4: 163-85.

Schoemaker, Paul and Howard Kunreuther. 1979. "An Experimental Study of Insurance Decisions." *Journal of Risk and Insurance* 46: 603

Ravallion, Martin and Michael Lokshin. 2000. "Identifying Welfare Effect from Subjective Questions." *Economica* 68: 335-57.

Scheve, Kenneth and David Stasavage. 2006. Stasavage and Scheve. 2006. "Religion and Preferences for Social Insurance", *Quarterly Journal of Political Science*, vol.1, no.3, pp.255-286.

Scott, J., Matland, P. Michelbach, and B. Bornstein. 2001. "Just Deserts: An Experimental Study of Distributive Justice Norms." *American Journal of Political Science* 45: 749-767.

Shoemaker, Paul and Howard Kunreuther. 1979. "An Experimental Study of Insurance Decisions." *Journal of Risk and Insurance* 46: 603-18.

Slovic, Paul, Baruch Fischhoff, Sarah Lichtenstein, Bernard Corrigan and Barbara Combs. 1977. "Preference for Insuring against Probable Small Losses: Insurance Implications". *The Journal of Risk and Insurance* 44: 237-258.

Starmer, Chris. 2000. "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk." *Journal of Economic Literature* 38: 332-382.

Varian, H. 1980. "Redistributive taxation as social insurance", *Journal of Public Economics* 14, 49-68

Figure 1: Tax Rates and Earnings in the Redistribution Game




Figure 1b: Tax Rates and the Difference Between Subject and Mean Earnings in the Redistribution Game

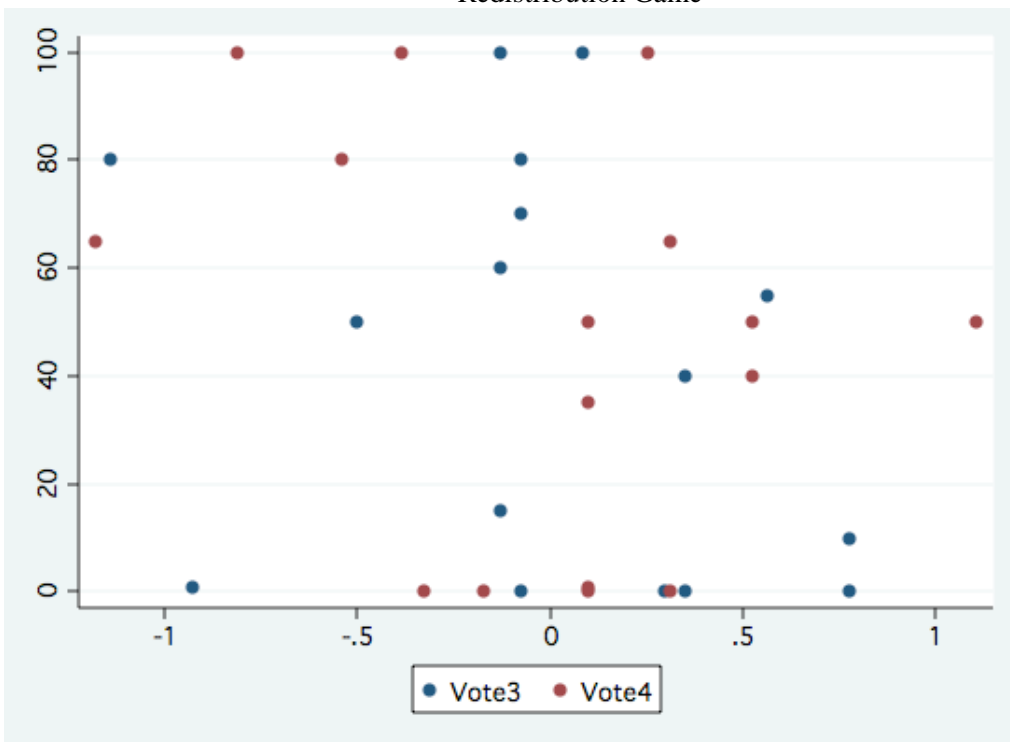



Figure 2: Tax Rates for Participants Below and Above Mean Earnings in the Redistribution Game



“Low” refers to earnings below the mean. “High” refers to earnings above the mean.

Table 1: Determinants of Tax Rates under different economic and institutional scenarios

WAVE 1	<i>No risk. Benefits allocated per capita</i>	
	Model 1: β(s.e.)	Model 2: β(s.e.)
Earnings in the previous period	-9.72 (10.73)	-1.82(9.6)
Round dummy	4.7(6.01)	7.4(6.2)
Ideology	--	6.16(5.68)
Cons.	50.33(20.9)	19.14(2.14)
N/Adj.R.sq.	50/0.03	40/0.06
WAVE II	<i>Insurance. Targeted benefits</i>	
	Model 1: β(s.e.)	Model 2: β(s.e.)
Earnings in the previous period	-7.010 (3.17)	-.972 (4.11)
Risk	100.1(26)	151.02(36.9)
Earnings*Risk	----	-34.6(19.2)
Risk Aversion	-23.01 (20.89)	-24.4(20.9)
Ideology	7.32(4.9)	5.84(4.85)
Round	7.69(7.6)	7.91(7.1)
Constant	6.43(11.4)	1.67(11.27)
N/Adj.R.sq.	40/0.451	40/0.478
WAVE III	<i>Insurance. Universal benefits combining insurance and redistribution</i>	
	Model 1: β(s.e.)	Model 2: β(s.e.)
Earnings in the previous period	-4.49(5.25)	-24.11(6.69)
Risk	144.39(38.84)	67.23(38.2)
Earnings*Risk	---	82.13(25.9)
Risk Aversion	15.70(31.2)	27.4(26.08)
Ideology	11.28(5.03)	9.79(4.7)
Round	-9.3(6.55)	-15.8(0.06)
Constant	-13.49(21.6)	13.59(16.8)
N/Adj.R.sq.	40/0.490	40/0.574

Model 1: Additive Specification. Model 2: Analysis of the interaction between pre-tax earnings in the previous period and risk exposure. All models are OLS with robust standard errors clustered by participant.

Table 2: Predicted Tax Rates: Conditional Effects in Wave II vs. Wave III

WAVE II: Insurance. Targeted Benefits					
<i>Pre-tax Earnings in previous round</i>					
<i>Risk</i>	-1.06	1.07	1.7	2.13	2.98
0	-- (*)	--	--	--	--
.1	26.3	16.8	14.08	12.18	--
.2	45.12	28.28	23.30	19.90	13.18
.35	73.29	45.37	37.11	31.48	20.34
.5	100	62.46	50.92	43.05	27.49
WAVE III: Insurance. Universal benefits mixing insurance and redistribution					
<i>Pre-tax Earnings in previous round</i>					
<i>Risk</i>	-1.06	1.07	1.7	2.13	2.98
0	52.24	--	--	--	--
.1	50.25	16.39	--	--	--
.2	48.27	31.9	27.06	23.7	17.2
.35	45.30	55.17	58.09	60.08	64
.5	43.13	78.44	89.12	96.4	100

(*) – reflects the fact that the 95% confidence interval of a particular prediction includes 0.
 The specific values for the IVs correspond, respectively, to the percentiles 1,25,50, 75 and 99

