DISCUSSION PAPER SERIES

IZA DP No. 5699

Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment

Guillermo Cruces
Ricardo Pérez Truglia
Martin Tetaz

May 2011

Forschungsinstitut
zur Zukunft der Arbeit
Institute for the Study
of Labor

Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment

Guillermo Cruces
CEDLAS-UNLP, CONICET and IZA

Ricardo Pérez Truglia
Harvard University

Martin Tetaz
CEDLAS-UNLP

Discussion Paper No. 5699
May 2011

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

IZA Discussion Paper No. 5699
May 2011

ABSTRACT

Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment*

Individual perceptions of income distribution play a vital role in political economy and public finance models, yet there is little evidence regarding their origins or accuracy. This study examines how individuals form these perceptions and posits that systematic biases arise from the extrapolation of information extracted from reference groups. A tailored household survey provides original evidence on the significant biases in individuals' evaluations of their own relative position in the distribution. Furthermore, the data supports the hypothesis that the selection process into the reference groups is the source of those biases. Finally, this study also assesses the practical relevance of these biases by examining their impact on attitudes towards redistributive policies. An experimental design incorporated into the survey provides consistent information on the own ranking within the income distribution to a randomly selected group of respondents. Confronting agents' biased perceptions with this information has a significant effect on their stated preferences for redistribution. Those who had overestimated their relative position and thought of themselves relatively richer than they were demand higher levels of redistribution when informed of their true ranking. This relationship between biased perceptions and political attitudes provides an alternative explanation for the relatively low degree of redistribution observed in modern democracies.

JEL Classification: D31, D83, H24, H53, I30

Keywords: perceptions of income distribution, limited information, preferences for redistribution, field experiment

Corresponding author:

Guillermo Cruces
CEDLAS, Facultad de Ciencias Económicas
Universidad Nacional de La Plata
Calle 6 entre 47 y 48, 5to.
Piso oficina 516 (1900) La Plata
Argentina
E-mail: gcruces@cedlas.org

* Funding for the *Encuesta de Percepciones Distributivas y Redistribución* was generously provided by Spain's Fundación Carolina (www.fundacioncarolina.es/es-ES/cealci/) and by PEGNet – the Poverty Reduction, Equity and Growth Network (www.pegnet.ifw-kiel.de). Additional funding was made available in the form of SECYT (PICT) and CONICET (PIP) grants. Part of this work was developed while Guillermo Cruces was visiting STICERD (LSE) in January and February 2008, with the support of a British Academy Visiting Fellowship. The opinions in this paper are those of the authors and do not represent the views of the funding agencies or of the institutions to which they belong. The authors are indebted to Raj Chetty, who provided detailed comments on an early draft of this paper, and to Alberto Alesina for valuable feedback on the first completed version. The authors also wish to thank François Bourguignon, Frank Cowell, Rafael Di Tella, Gary Fields, Leonardo Gasparini, Andrés Ham, Matt Rabin, Walter Sosa Escudero, Julián Troksberg and Carlos Winograd for valuable discussions on the project leading to the preparation of this paper. They are also grateful for comments made by Alejandro Bonvecchi, Juan Dubra, Emilio Espino, Chico Ferreira, Laura Guardia, Juan Carlos Hallak, Alejandro Manelli, Andrés Neumeyer, Hugo Ñopo, Laura Ripani, Martín Rossi, Ernesto Schargrodsky, Jacques Silber, Mariano Tommasi and other seminar participants at the 2009 PEGNet conference (The Hague, September 2009), LACEA (Buenos Aires, October 2009), Universidad de San Andrés (2010), NIP (Universidad de la República, Montevideo, 2010) and Universidad Torcuato Di Tella (Business School, 2010, and Economics Department, 2011). Darío Tortarolo and Florencia Pinto provided a meticulous consistency analysis of the survey data. The usual disclaimer applies.

1 Introduction

The shape of the income distribution plays a key role in the determination of policies with redistributive components (such as social security, health care, government transfers and taxation) in political economy and public finance models. However, the main policy determinant is not its actual shape, but rather how it is perceived by agents in the economy. This study fills a gap in the literature by exploring the causes and consequences of systematic biases in individuals' perceptions of aggregate income distributions.

The findings presented in this paper contribute to the recent literature on the incorporation of subjective perceptions and inference problems into the determination of political economy outcomes (for a seminal contribution see Piketty, 1995). For instance, when forming their views on public policies, agents may need to infer the importance of effort and predetermined factors in the income generation process (Piketty, 1995; Bénabou and Tirole, 2006), they may evaluate prospects of economic mobility (Bénabou and Ok, 2001; Alesina and La Ferrara, 2005), or they may arrive at conclusions as to the causes of poverty and the fairness of socioeconomic outcomes in general (Alesina and Glaeser, 2004). To form their judgments, views and attitudes, agents need to make difficult inferences about distributional outcomes (e.g., inequality, mobility) based on limited information and within given time constraints, but there is as yet little evidence on the origins or the accuracy of the inferences they make in this regard.

This paper also makes a contribution to a growing body of work that attempts to document agents' expectations and subjective probabilities (Manski, 2004; Hurd, 2009) and explain their formation (Zafar, 2011). In an application to distributional issues, Norton and Ariely (2011) elicit information on American's perceptions of the wealth distribution in their society and find significant discrepancies between actual and perceived levels of inequality. The results presented in this paper not only document systematic discrepancies between objective and subjective income distributions, but also provide a series of tests for the origins and the consequences of such discrepancies.

The assessment of an income distribution by an economic agent can be regarded, fundamentally, as a statistical inference problem. Individuals observe the income levels of no more than a sub-sample of the population and must infer the entire distribution from that information. If agents do not fully account for the selection process involved in the formation of the sample they observe, their inferences will be systematically biased. This failure may be due to limitations in the information set available to the agents – the information might be

costly or unavailable. Alternatively, agents may have the necessary information, but they may sometimes fail to use it correctly, as argued in the cognitive bias literature (Rabin, 1998; Camerer, Loewenstein, and Rabin, 2003). Irrespective of whether agents have limited information or bounded rationality, this rationalization of distributional perceptions provides a series of corollaries that can be tested with data on objective and perceived distributions.

The empirical results in this paper are based on the *Survey on Distributional Perceptions and Redistribution*, a study of 1,100 households representative of the Greater Buenos Aires area in Argentina. The survey was designed and implemented in 2009 for the specific purpose of testing the posited mechanisms for the formation of distributional perceptions. Data were collected on each respondent's household income and on his or her assessment of its ranking (to the closest decile) in the overall income distribution.

The first finding is the presence of systematic biases in perceptions of own income rank: a significant portion of poorer individuals place themselves in higher positions than they are, while a significant proportion of richer individuals underestimate their rank. Moreover, as predicted by the theory, the bias is significantly correlated with the respondent's relative position within the reference group (as proxied by area of residence). Also in keeping with the theory, respondents with friends from heterogeneous social backgrounds are less prone to these biases.

Finally, the analysis explores how these misperceptions about the income distribution may affect attitudes towards redistribution. For instance, self-interest might induce poor individuals to demand less redistribution if they think they are relatively richer than they actually are. The document presents the results from a unique randomized experiment that was implemented within the survey: for a randomly assigned treatment group, the interviewer highlighted any discrepancy between the subjective assessment of the respondent's ranking and that respondent's actual position, effectively correcting any bias that was present. This survey field experiment contributes to the literature on information provision as a treatment (Duflo and Saez, 2003; Chetty and Saez, 2009; Card et al., 2010). An original feature is that perceptions are not only contrasted with reality (as in Olken, 2009, among others): the survey experiment provides feedback and actually confronts biased subjects with consistent information.

The results from the experiment indicate that confronting agents' biased perceptions with consistent information had a significant effect on stated preferences for redistribution. Those who overestimated their relative position (who thought of themselves to be relatively richer than they were) and were provided consistent information demanded more

redistribution than those in the control group. To the degree that the information treatment managed to correct biased distributional perceptions, these results can be interpreted as evidence of the effect of biases in distributional perceptions on political attitudes. Specifically, upward biases in perceptions of own rank among the less well-off reduce their demands for redistribution. This finding constitutes an alternative to theories that posit prospects of upward mobility (Bénabou and Ok, 2001) or other factors as accounting for the relatively low levels of demand for redistribution in modern democracies.

This paper is organized as follows. The next section discusses the formation of subjective income distributions and individuals' perceptions of their income rank, and explores these factors' implications for attitudes towards redistribution. The third section describes the household survey and the randomized experiment that was designed to answer these questions. The fourth section presents the empirical results. The last section concludes.

2 Subjective income distributions and potential biases

2.1 Subjective income distributions and reference groups

Economic agents' assessments of income distributions depend on their access to information and on their ability to process the relevant data. The latter constitutes a trivial matter in a perfect information context, where the incomes of all members of society are observed. However, in the presence of limited information, these assessments become statistical inference problems. Individuals are constantly exposed to the income levels of others through, for instance, the media and social interaction with acquaintances, co-workers, employees, etc. Even if they interact primarily with individuals exhibiting similar characteristics, rational agents may arrive at consistent estimates of the entire distribution by factoring in the selection process of the non-representative sample of incomes that they observe.

The distribution of income in society as a whole is given by the density function $f(\cdot)$. An agent *j* can infer $f(\cdot)$ (or some statistic of $f(\cdot)$ – the mean, the median, or the agent's own ranking in the distribution) using the information about the incomes that he or she has observed. This sub-sample of observed incomes constitutes agent *j*'s reference group, S_j. The term $\mathrm{P}(\mathrm{i} \in S_j)$ denotes the probability that a randomly chosen individual *i* belongs to group S_j, and P_{S_j} denotes the proportion of the whole population that belongs to S_j.

Let $P(i \in S_j \mid x_i)$ be the selection equation, which represents the probability that individual *i* belongs to the reference group S_j given that his or her income is x_i. From the observed incomes, the agent can derive a consistent estimate of the distribution of incomes within her or his reference group, $f(x_i \mid i \in S_j)$. That density distribution is related to the unconditional density distribution $f(x_i)$ (i.e., the density of the entire distribution of incomes, not just group *j*) by the following identity:

$$f(x_i \mid i \in S_j) = \frac{P(i \in S_j \mid x_i)}{P_{S_j}} f(x_i) \qquad (1)$$

Agent *j* is deemed sophisticated (subscript S) if he or she applies Bayes' rule to infer the income distribution for the entire population from the observed subset (S_j):

$$f_j^S(x_i) = f(x_i \mid i \in S_j) \frac{P_{S_j}}{P(i \in S_j \mid x_i)} \qquad (2)$$

To make the inference given by equation (2), agent *j* requires information about the relative size of the reference group (P_{S_j}), knowledge about the selection process leading to the formation of that reference group ($P(i \in S_j \mid x_i)$), and the ability to make probability judgments. Any deviation from these conditions will result in biased perceptions about $f(x_i)$.

There are several sources of deviations from these conditions. For instance, the information about $P(i \in S_j \mid x_i)$ or $f(x_i)$ may be costly to acquire, or the advantages of doing so might not be evident. In equilibrium, they may opt for good enough naïve estimates. It may also be the case that, as in Benoit and Dubra (2011), the naïve estimate represents the best possible answer that can arise from rational agents' extrapolations conditioned on the information set available to them. The key factor in this situation is the limited information about reference groups or about the selection process leading to reference groups.[1] Alternatively, individuals may fail to consider all the available information, or they may use it incorrectly (Simon, 1972). For example, agents may use heuristics or rules of thumb when dealing with difficult questions of statistical inference, and such rules of thumb can, on occasion, be very imprecise. Indeed, the use of heuristics in statistical inference and the systematic biases that such exercise entails is a well-documented phenomenon in the

[1] A special case is given by a complete lack of information about a subgroup of the population. If $P(i \in S_j \mid x_i)$ is zero for some x_i (i.e., the individual does not observe some income ranges with probability one), then even with the ability to make probability judgments and with full knowledge of $P(i \in S_j \mid x_i)$ the agent will still not be able to apply Bayes' rule.

cognitive literature (Kahneman, Slovic and Tversky, 1982). The most relevant case in this discussion is the representativeness heuristic, in which individuals fail to apply Bayes' rule to the information they obtain (Kahneman and Tversky, 1972; Tversky and Kahneman, 1974). This failure leads to a systematic cognitive bias, the base rate neglect.[2]

These two possibilities, limitations in information and bounded rationality, can be illustrated by a situation in which a "naïve" agent does not properly apply Bayes' rule as in equation (2). In the extreme case, irrespective of the source of the bias, the naïve agent will simply use the information about the income distribution within his or her reference group as if it were representative of the entire population. The beliefs of naïve agents in this extreme case are denoted by the superscript N: $f_j^N(x_i)$. Equation (3) establishes the relationship between the perceptions of sophisticated and naïve agents:

$$f_j^N(x_i) \equiv f(x_i \mid i \in S_j) = f_j^S(x_i)\frac{P(i \in S_j \mid x_i)}{P_{S_j}} \qquad (3)$$

From $f_j^N(x_i)$ it is straightforward to obtain the naïve estimates of many characteristics of the income distribution, like the mean, median, dispersion, and proportion of individuals under the poverty line, $F_j^N(x_{poor})$, among others. Therefore, any biases in $f_j^N(x_i)$ are likely to generate biases in a wide set of perceptions of the income distribution.

If the formation of reference groups is orthogonal to income, then the selection equation will be such that $P(i \in S_j \mid x_i) = P(i \in S_j)$; that is, on expectation, every group will be representative of the whole population. In this case, it would be consistent to use the within-group distribution as an estimate of the distribution for the entire population. The selection equation, however, is probably a function of income. A probability $P(i \in S_j \mid x_i)$ increasing in x_i represents the case of an agent *j* who has a "rich" reference group and is therefore more likely to observe higher-income individuals. Conversely, if agent *j* interacts mainly with lower-income individuals (i.e., a "poor" reference group), then $P(i \in S_j \mid x_i)$ is decreasing in x_i.

This is illustrated in Figure 1. Figure 1a depicts the income distribution for the whole population and for a rich reference group, which exhibits fist-order stochastic dominance over the distribution for the whole population (i.e., for every income level in the reference group there is a greater share of people below that income level than in the whole

[2] The base rate neglect has been incorporated in economic models and empirical applications before (see Grether, 1980, 1990; Rabin, 1998; Camerer et al., 2003; DellaVigna, 2009).

population). Since naïve agents use the information about the income distribution within their reference group as if it were representative of the entire population, naïve agents in the rich reference group will underestimate the actual cumulative income distribution for every income level, i.e., $F^N(y) < F(y)$ for all y. In Figure 1a, this is illustrated for a given income y_1 by the difference between the areas filled with horizontal and vertical lines. Conversely, naïve agents with poor reference groups will overestimate the cumulative income distribution for every income y. The results are not straightforward when there is no stochastic dominance of the distribution within a reference group over that of the whole population. Figure 1b illustrates this result by showing a "middle-class" reference group, where agents underestimate F(y) for income levels y<y* and overestimate F(y) for incomes y>y*.

There is an intuitive way to relate the size and direction of the bias in $F_j^N(y)$ to the selection process for reference groups, $P(i \in S_j \mid x_i)$. Taking the integral of equation (3) from 0 to y, some algebraic manipulation results in the following expression:

$$F_j^N(y) = \int_0^y f_j^N(x_i)dx_i = F_j^S(y)\frac{\int_0^y P(i \in S_j \mid x_i)\frac{f_j^S(x_i)}{F_j^S(y)}dx_i}{P_{S_j}} = F_j^S(y)\frac{R_j(y)}{P_{S_j}} \qquad (4)$$

The term $R_j(y)$ represents the average probability of belonging to group S_j for individuals with incomes below y. If agent *j* has a rich reference group, those with incomes lower than y have a lower probability of belonging to the rich reference group than the average of this probability for the entire population, so $R_j(y) < P_{S_j}$. As a result, agents with rich reference groups underestimate the entire cumulative distribution function: $F_j^N(y) < F(y)$ for all y.

The expression $f_j^N(x_i)$ in equation (3) can be used to compute naïve perceptions of different moments and statistics of the income distribution. For instance, if reference groups are more homogeneous in income than the whole population (as it is likely to be the case), perceptions about income inequality will be biased downwards for all agents. Indeed, this is consistent with Norton and Ariely's (2011) finding that individuals systematically under-estimate the level of inequality in the distribution of wealth in the United States.

A crucial parameter for this study is $F_j^N(x_j)$, the perception of an agent's own income rank within the distribution. Since agents with rich reference groups underestimate all points in the cumulative distribution (including x_j), it follows that they will underestimate their own

rank in the distribution. Conversely, naïve agents with poor reference groups will over-estimate their rank.

Finally, it is important to note that these naïve estimates represent extreme cases in which agents completely fail to take into account the selection process of their reference group. The model can be generalized by letting individuals make imperfect inferences using some information about the selection equation, $P(i \in S_j \mid x_i)$. If the individual under-estimates the importance of income in the formation of reference groups, then the biases will be qualitatively similar to those presented in the extreme naïve case.

2.2 Biased perceptions and preferences for redistribution

Misperceptions of income distribution can have substantial implications in the determination of policy outcomes. This can be illustrated by incorporating biased perceptions into a basic Meltzer and Richard (1981) type of framework with a simple redistributive policy reform in which taxes would finance some benefits. The population can be divided by income levels into potential net losers and winners: individuals above a given π-quantile would lose in net terms if the redistributive policy were implemented, while those below it would gain from the reform. For instance, with a linear income tax and lump-sum redistribution, agents below the median income will benefit and those above the median will lose.

If agents have biased perceptions of their own income ranking, their evaluations of how these costs and benefits will affect them are likely to be inaccurate. The agents that will benefit from the policy are those whose actual relative income is below the π-quantile: $F(y_j) < \pi$. Since they are relatively poor, most of these agents are likely to have poor reference groups, and naïve agents among them will therefore over-estimate their own ranking. Some of these naïve agents will consider themselves potential non beneficiaries of the reform, since their perceived relative income is above the π-quantile: $F_j^N(y_j) > \pi$. Thus, there will be a group of agents who erroneously believe that they would not benefit from further income redistribution when they actually would: i.e., those with incomes between $y_\pi^L = F^{-1}(\pi)$ and $y_\pi^H = \left(F_j^N\right)^{-1}(\pi)$. With access to the correct information about their actual income rank, self-interest would make these individuals change their attitude and favor, rather than oppose, the redistributive policy. Analogously, a group of naïve individuals with rich reference groups will believe that they would benefit from the redistributive policy,

although they would actually be net contributors. If those individuals are purely self-interested, providing them with consistent information about their income ranking would make them oppose, rather than favor, the redistributive policy. The experimental design of the survey used in this study allows for a direct test of this mechanism.

However, this simple framework has some shortcomings. The recent political economy literature has discussed at length the fact that individuals may be motivated by more than self-interest, so preferences for redistribution may reflect a wider set of factors. For instance, agents may have altruistic motives and incorporate fairness considerations into their decisions regarding their support for redistribution.[3] As discussed above, the biases given by equation (3) translate into biased perceptions of many characteristics of the distribution of income, including the poverty count and social welfare, which could influence an individual's support towards redistributive policies through altruistic concerns. The simplest illustration is that of aggregate deprivation. With a poverty line z, $F(z)$ is the actual poverty headcount and $F_j^N(z)$ its biased perceived level. Naïve individuals with rich (poor) reference groups will underestimate (overestimate) the aggregate poverty headcount.[4] If preferences for redistribution are increasing in the perceived poverty level, then altruistic naïve agents with rich (poor) reference groups would support more (less) redistribution upon correcting their biased perceptions. As a result, providing consistent information to naïve individuals may have conflicting effects in terms of their support for redistribution from the perspective of selfish and altruistic motives. The experiment described below, however, identifies only the net effect of providing consistent information on the income distribution on attitudes towards redistributive policies.

[3] Fong (2001), Luttmer (2001), Rottemberg (2002), Alesina and Glaeser (2004) and Alesina and Angeletos (2005) study the effect of altruistic and fairness concerns on attitudes towards redistribution. See Alesina and Giuliano (2009) for more references.

[4] This conclusion is part of a more general result. For any social welfare function increasing in income (Cowell, 2000), first-order stochastic dominance of an income distribution over another implies that social welfare under the first distribution is greater than social welfare under the second distribution. If the distribution of incomes in the rich (poor) neighborhood dominates (is dominated by) the distribution in the whole population, it follows that naïve individuals in the rich (poor) reference groups will overestimate (underestimate) true social welfare.

3 Data source and experimental setup: Survey on Distributional Perceptions and Redistribution

3.1 Survey on Distributional Perceptions and Redistribution

The discussion in the previous section covered the formation of subjective income distributions, the possibility of systematic biases, and their implications for attitudes towards redistribution. The empirical investigation in this paper is based on the *Survey on Distributional Perceptions and Redistribution*, a study of 1,100 households representative of Greater Buenos Aires in Argentina. The survey was carried out in March 2009 and consisted of face-to-face interviews of a random sample of that population. It was specifically designed to test the model presented in the previous section and, to that end, collected data on a set of individual and household characteristics and on respondents' labor-market and other socioeconomic outcomes, as well as their answers to a series of questions about their political views and attitudes. It also gathered information on the respondents' actual household income and on their perceptions of their own income rank within the distribution for the whole country.

There are several ways of recovering subjective probability distributions for a continuous variable such as income, which include eliciting quantiles, moments or points of the distribution (see Manski, 2004). For instance, Norton and Ariely (2011) collected information on respondents' evaluations of the proportional distribution of total wealth among quintiles in the United States. The *Survey on Distributional Perceptions and Redistribution* relied on an original instrument (the income-rank evaluation question), which elicited a specific value of the cumulative subjective distribution: its evaluation at the point where each respondent thought his or her household stood.[5] The question was worded as follows: "*There are 10 million families in Argentina. Of those 10 million, how many do you think have an income lower than yours?*"[6] The survey also collected data on the households' total monthly income by intervals. While distributional indicators often rely on per capita or adjusted income, a pilot conducted in December 2007 indicated that individuals compare incomes in terms of total household monthly levels. The intervals were chosen by the research team to correspond to the boundaries of deciles of national total household income

[5] Nuñez (2005) collected information about the respondents' evaluation of the percentage of households above and below their income level in Chile. The approximate number of households in Argentina at the time of the survey (10 million) allowed the question to be phrased in terms of millions of households on a simple 1-10 scale, thus eliminating the need for respondents to be comfortable in answering in percentage terms.

[6] This information differs conceptually from measures of subjective economic welfare (Ravallion and Lokshin, 2002), since it attempts to capture an objective parameter of the distribution. In this sense, it is closer to the literature on elicitation of subjective probabilities (Hurd, 2009).

distribution at the time of the survey to facilitate the comparison of objective and perceived position in the distribution in the experimental design.[7]

3.2 The survey experiment setup

Besides the income rank question, the second and most innovative aspect of the survey was the implementation of an experimental design incorporated into the questionnaire. Randomized questionnaire-experiments had been developed in laboratory settings (Amiel and Cowell, 1992; Cowell and Cruces, 2004), while, in the context of household surveys, Joliffe (2001) and Di Tella, Galiani and Schargrodsky (2008) have randomized the allocation of questionnaire types among respondents.[8]

As in these previous studies, the experimental setup for this survey involved randomly allocating two different types of questionnaires to interviewees, although the questions asked to the respondents were the same. The originality of this setup stems from the nature of the treatment, in which the interviewer provided feedback to respondents in the treatment group in the form of consistent information on the income distribution. Specifically, after collecting information on household characteristics, income levels and positional perceptions, the interviewer informed respondents in the treatment group whether their estimates of relative income coincided with those of the research team. The interviewer read the following statement (with X and Y being determined by previous answers): *"Based on your income level, the latest studies conducted by the University indicate that there are X million families with an income lower than yours, while you stated that there were Y."* The interviewer then read out one of the three following statements, depending on the accuracy of the X/Y comparison: (1) *"In fact, there are more families with a lower income than yours than you believed"*, (2) *"You were right about how many families have a lower income than yours"*, or (3) *"In fact, there are fewer families with a lower income than yours than you believed."* The presence of a bias in their perceptions was thus explicitly pointed out to respondents in the treatment group. After the treatment, the questionnaire was used to collect information on attitudes about specific redistributive policies of interest in Argentina within the political context existing at the time of the survey. The questionnaire for the control group did not

[7] The use of income intervals significantly reduces non-response rates. The notes in appendix table A1 provide further details on the construction of the intervals and their implementation by interviewers in the survey.

[8] Survey experiments have also been conducted in the context of political science and public opinion research (see, for instance, Horiuchi et al., 2007; Hainmueller and Hiscox, 2010).

contain the "feedback" section, but was exactly the same in all other respects (Table A1 presents an extract of the questionnaire and variable definitions).

This experimental survey design contributes to a growing body of literature concerning the provision of information as a treatment in an experimental setting. For example, Duflo and Saez (2003) and Chetty and Saez (2009) provided subjects with information on retirement plans and the tax code, respectively, while Jensen's (2010) study offered statistics on returns to schooling for teenagers, and Card et al. (2010) gave a group of employees access to information on co-workers' wages. There are also several studies that have contrasted subjective and objective probabilities and their relationship with actual outcomes in connection with, for instance, income expectations versus realizations (Manski, 2004), objective versus subjective income percentiles (Nuñez, 2005) and perceived versus actual survival rates (Hurd, 2009). This study innovated in a crucial way, however, by confronting subjects with consistent information which differed from their stated perceptions.

4 Empirical results and discussion

4.1 Subjective income distributions

This section presents the empirical evidence from the *Survey on Distributional Perceptions and Redistribution*. The starting point for the analysis is the distribution of objective and perceived income rank in the sample. Figure 2a presents the income distribution of the Greater Buenos Aires survey sample as a function of deciles of the national distribution at the time of the survey, which served as the basis for the categories used for the household income question. A nationally representative sample would be depicted in the figure as a horizontal line at the 10 percent density. The higher concentration in deciles 5 to 9 is accounted for by higher average income levels in Buenos Aires with respect to the country average. Figure 2b, in turn, presents the respondents' perceptions of their households' positions in the distribution, which were elicited by posing the income-rank evaluation question described in the previous section. By construction, respondents identified what decile of the national distribution they thought was the closest to their income level. In contrast with the fairly even distribution in Figure 2a, the mode of the perceptions distribution is given by the fifth decile, with almost 30 percent of respondents placing their households at that level (and almost half in the middle quintile – corresponding to the fifth and sixth deciles). Self-perceptions of income rank in the sample are thus substantially less dispersed than objective income levels are.

The difference between the two panels in Figure 2 indicates the presence of a bias in distributional perceptions. The bias is defined here as the difference between a household's objective income decile and the respondent's self-assessment of its position (in deciles): those with a negative bias consider themselves to be in a lower position than they really are, while the opposite is true for those with a positive bias. The distribution of this variable, depicted in Figure 3a, indicates that there are a significant number of cases of both positive and negative biases. In fact, only about 15 percent of the respondents place their household's income in the objectively correct decile. However, the deciles of a national income distribution are relatively narrow categories, and it is therefore quite plausible that respondents could have difficulty in ascertaining their position with that degree of precision. To allow for measurement error in agents' perceptions (Bertrand and Mullainathan, 2001) and in the levels of the objective decile boundaries, Figure 3b presents an alternative definition of this bias. Here, only respondents who deviate from their true position by two deciles or more are considered to be biased. As expected, the number of biased responses is substantially lower than in Figure 3a. However, Figure 3b still indicates that, even with this less demanding definition, more than 55 percent of the respondents exhibit some degree of bias, with relatively more cases of negative biases (individuals placing themselves below their true ranking) than positive ones.

A further question is whether there is a relationship between the distributions depicted in Figure 2. Table 1 and Figure 4 present the relationship between objective and perceived relative income levels. Figure 4a depicts the average of perceived own-income deciles, by levels of objective deciles. There is a significant positive relationship between both variables, although the distribution of subjective levels is considerably more concentrated. For instance, the average perceived own-income decile fluctuates around 4.5 for the three poorest deciles of the objective income distribution, but is less than 7 for those in the top objective decile. This pattern has a direct correlate for the distribution of the bias, which is depicted in Figure 4b as a function of the objective income decile: respondents at the top and the bottom of the objective distribution display substantial negative and positive biases, respectively (of about -3.5 and 3 deciles for the extreme categories). Moreover, the bias diminishes up to the fifth objective decile and increases monotonically (in absolute values) from the sixth decile onward. Table 1 presents similar information on the bias for a partition of the population in quintiles of the objective income distribution. The table also indicates that positive biases are largely confined to respondents below the median of the distribution, while those with a negative bias are concentrated in the fourth and fifth quintiles.

4.2 Reference groups and biased perceptions of income distribution

The discussion so far has revealed the presence of substantial biases in distributional perceptions. Section 2 posited a mechanism for the formation of subjective income distributions, whereby individuals extrapolated from information about the income distribution in their reference groups in order to obtain estimates of the whole distribution. Sophisticated agents (rational and informed) can obtain a consistent estimate of the aggregate distribution even if their reference group is not fully representative by correctly applying Bayes' rule (that is, accounting for the selection process of the sample that they are observing). By contrast, a "naïve" individual fails to apply this rule and considers only his or her reference group when making inferences about the whole population; this results in systematic biases if selection into reference groups is a function of income.

If reference groups bundle together individuals of similar income levels, then one simple prediction of the model is that individuals with rich reference groups (and therefore most rich individuals) tend to underestimate their income rank, whereas individuals with poor reference groups (and therefore most poor individuals) overestimate their rank. This distribution of biases is exactly the one depicted in the results above, most notably in Figure 4b and Table 1. However, the observed pattern is also consistent with other potential explanations. For instance, individuals may have a tendency to the mean (or the median), as has actually been documented in the literature on expectations and subjective probabilities,[9] such that the middle deciles constitute focal-point answers to the income-rank question. Alternatively, poorer respondents may feel embarrassed to admit that their relative income is low and thus over-report their true (accurate) perception, while richer individuals may not feel comfortable reporting their high relative position[10] and thus under-report their true (accurate) perception. The survey allows for the use of a series of empirical means of testing the reference-group hypothesis against these simple mechanical explanations.

In the discussion presented in Section 2, the entire set of individual interactions (with friends, family, co-workers, etc.) was regarded as constituting the relevant reference group for the formation of perceptions of income distribution. The analysis here uses a geographical

[9] Hurd (2009) points out that "when the true probability of an event is greater than 0.50 […] the subjective probability will be understated" and vice versa, and finds several examples of survey responses with focal points at 50 for distributions between 0 and 100. Similar factors might be at work in the evaluation of income distributions. For instance, 50.4 percent of respondents in the 1972-2008 pooled General Social Surveys in the United States stated that their income was "about average" (own tabulations).

[10] Under-reporting of income for higher levels is typically a concern in household surveys. However, in this case, the tendency of those with higher income levels to underestimate their position implies that under-reporting at the top of the distribution would reduce the number of those classified as biased. The substantial number of respondents with a negative bias can be considered to be a lower bound.

approximation: an individual's reference group is given by his or her area of residence. This approximation is justified by the significance of social networks at the neighborhood level for the exchange of information on employment and other income-generating activities (Topa, 2001; Bayer et al., 2008).[11] Moreover, areas of residence provide a simple illustration for a reference-group selection mechanism based on income levels, given the pervasive residential segregation of households by income levels in urban areas (Glaeser et al., 2008).

The survey covered 41 randomly selected sampling points within 10 localities of the Greater Buenos Aires metropolitan area. The sampling points correspond to a fairly small set of street blocks and contain 26 households on average – these are referred to as "neighborhoods" in the discussion. The average objective income level reported in the survey within each neighborhood is depicted along the horizontal axis of Figure 5a. This figure reveals substantial variation in average incomes for neighborhoods in the sample, which range from an average objective decile of 3 to just below 8.

The hypothesis that a salient role is played by geographic reference groups in the formation of distributional perceptions (and misperceptions) predicts a strong relationship between the economic level of the residence area (as a proxy for the reference group) and the direction of the biases. Figure 5a illustrates the strong positive correlation between average income levels and average perceptions at the neighborhood level.[12] Figure 5b, in turn, fits the prediction of the model discussed in Section 2: individuals in lower-income neighborhoods (i.e., with average objective income ranks below 4.5) overestimate, on average, their income rank (positive average bias), and individuals in higher-income neighborhoods (i.e., with average income ranks above 6.5) underestimate, on average, their income rank (negative average bias). The model also predicts that agents in "middle-class" neighborhoods will have a balanced mix of positive and negative biases, possibly generating a zero average bias. This is indeed the case for neighborhoods with average income deciles between 4.5 and 6.5.

This evidence suggests a role for geographic reference groups in the determination of subjective income distributions, since the above predictions are specific to the reference-group hypothesis.[13] The information collected in the survey allows for a sharper test. The discussion in Section 2 pointed out that fully naïve agents will report their positions within

[11] The income comparisons literature, which focuses mainly on the impact of relative income on subjective well-being, also employs geographical definitions of reference groups (see the review of Clark et al., 2008).

[12] The correlations in both panels of Figure 5 also hold by locality. Where possible, indicators aggregated at the geographical level are based on neighborhoods rather than localities to provide greater variation in the figures.

[13] For instance, if respondents to the survey simply provided a focal-point answer (5 plus/minus some random term) in reply to the question on the perceived own-income rank, there should be no relationship between average income and average perception by geographical area.

their reference groups as their perceptions of their income ranks within the whole population. On the other hand, under the null hypothesis that individuals correctly apply Bayes' rule, relative income within a reference group should not have any explanatory power for perceived income rank within the population as a whole after controlling for the agent's objective income rank. More specifically, under the alternative hypothesis that individuals under-estimate the importance of income in the formation of reference groups, then the relative income within the reference group will be positively correlated to their perceptions of income ranking.[14]

Table 2 presents a series of tests of this hypothesis that employ regression analysis. In all regressions, the perceived own-income decile is the dependent variable.[15] Column 1 presents a simple regression with the respondents' objective income deciles as the sole independent variable. The results in this column confirm the strongly significant relationship between the two variables discussed above. The level of the coefficient, however, reflects the differences in values between the two variables. This can be seen in Figure 4a: in a perfect information context, the objective income variable should have a coefficient of 1 and the regression should have an r-square of 1. The estimate is strongly significant, but the coefficient is only 0.2452. The second column repeats this simple regression, but includes 41 neighborhood fixed effects:[16] the coefficient is slightly lower, at about 0.201, and still significant at the 1 percent level.

The first direct test of the geographic reference group hypothesis is presented in column 3. The regression includes the respondents' income rank within their localities (the number of households within sampling points is too small to provide a meaningful measure), transformed to the 1-10 scale of the objective and subjective income decile variables. With the inclusion of the locality rank variable as an independent variable, the coefficient of the objective income decile variable is virtually zero, and not significant at standard levels. The coefficient of the locality rank variable, on the other hand, has a positive and statistically significant effect of 0.2151 on the perceived own-income decile. The respondents' relative incomes within their localities thus seem to have an effect on their perceptions of the distribution, over and above that of their own objective income levels. The coefficient is

[14] However, in principle it could also be the case that the individual over-estimates the role of income in forming reference groups. If that was the case, then the relative income within reference groups would be negatively correlated to the perceptions of income ranking.

[15] Regressions with the bias as the dependent variable do not convey meaningful results because, by construction, the bias is strongly correlated with the objective income decile.

[16] Standard errors are clustered at the neighborhood level. All the results in the table are maintained if 10 locality fixed effects (with robust standard errors) are included instead.

economically significant and is close to that of the objective income level in the simple regressions of columns 1 and 2. This result does not arise from a high collinearity between objective income and rank within locality, as shown by the regression in column 4, which includes objective income deciles as a series of 9 dummy variables. Even with this semi-parametric control for objective income, the coefficient of relative income within a locality is statistically significant and about the same in size as in column 3, and the F-test rejects the joint significance of the objective income dummies with a p-value of 0.24.

The following column in Table 2 presents the results of another robustness check. A potential concern may be that the measure of objective income is imprecise, so that the locality income rank variable may be indirectly capturing the effects of unobserved variations in actual income levels. The model in column 5 includes a set of additional regressors consisting of individual and household characteristics to proxy for the respondent's income-generating capacity (his or her education level, that of his or her spouse, age, gender, type of employment – see the Table's notes for details). If the locality rank captures some of the omitted variables, its coefficient should decrease substantially with the introduction of these controls. The results shown in column 5 of Table 2 indicate that adding this exhaustive set of controls does not significantly alter the point estimate or the statistical significance of the coefficient of the rank within locality variable, and these results are robust to the inclusion of further controls.

Finally, Table 2 also presents the results of a further test of the reference group hypothesis. The survey included a question intended to measure the breadth of respondents' reference groups: *"Among your friends and co-workers, would you say that there are individuals from all social classes (1), or, if not, that most of your friends belong to the lower class (2), the middle class (3), or the upper class (4)?"* The response is used to generate a dummy variable that takes the value 1 if an individual responded that his or her friends are from all social classes (38.5 percent of the respondents), and zero otherwise. Intuitively, people that interact in several reference groups must have more information about the role of income in forming reference groups (or alternatively, the selection process is more salient to them, so they are less likely to fail to consider Bayes' rule). As a consequence, they should be less prone to report their relative position within their locality as an estimate of their ranking in the entire distribution. The regression shown in column 6 includes the indicator variable for having friends from all social classes, as well as an interaction of this indicator with the respondent's income rank within the locality. The coefficient of this interaction should be negative: conditional on their own objective income level, these individuals should be less

influenced by the relative income within their reference group.[17] The results obtained from this test also support the reference group hypothesis: the coefficient of the interaction term is negative and both statistically and economically significant.

The empirical results shown in Sections 4.1 and 4.2 are consistent with the factors highlighted in the discussion in Section 2. Systematic biases in subjective income distributions are manifested here as deviations in the perception of own-income rank. A significant percentage of poorer individuals place themselves in higher positions, while a significant proportion of richer individuals think that they are closer to the median than they really are. Moreover, the results confirm the salient role played by reference groups (proxied by area of residence) in the formation of subjective income distributions. As predicted, there is a positive relationship between an individual's perception of income ranking and relative position within the individual's area of residence, over and above the effect of the individual's objective income level. Also in line with the predictions, individuals with friends from heterogeneous social backgrounds are less prone to be biased by their relative rank in their area of residence.

4.3 Biased perceptions and preferences for redistribution: experimental results

The empirical evidence presented so far illustrates the degree of bias in distributional perceptions, and the following analysis assesses their relevance for attitudes towards redistribution. Studying the effect of distributional misperceptions on political attitudes poses a significant identification challenge, since preferences for redistribution are significantly correlated with income levels (Alesina and Giuliano, 2009), which are in turn closely linked to an individual's perceived position and the resulting bias, as illustrated by the results in Table 1 and Figure 4. A simple analysis of political attitudes by objective and perceived level of income would thus be marred by this spurious correlation problem. This section presents an original empirical test to determine whether these misperceptions may affect economic decision-making. The test was designed to overcome these identification difficulties and to illustrate how misperceptions can have first-order implications for economic and political processes through their impact on individuals' attitudes towards redistribution and related policies.

[17] In an extreme case, if this variable indicates that the respondent has full information on the selection process, then the coefficient of the interaction should be negative and equal in absolute value to that of the geographical relative rank.

As described in Section 3, the survey included a field questionnaire-experiment: the interviewer provided a randomly assigned group of respondents with "correct" (i.e., unbiased) estimates of their positions in the income distribution based on their responses. The interviewer explicitly pointed out the degree and direction of the bias in each respondent's self-assessment (if any). The experiment thus identifies the causal effect of this information treatment on preferences for redistribution. Moreover, to the degree that the treatment manages to correct biased distributional perceptions,[18] the results can be interpreted as capturing evidence of the effect of misperceptions on political attitudes.

The focus in this study is on attitudes towards a specific set of policies rather than on general beliefs about justice and income redistribution. The policies under study consist of direct government transfers to the poor. These policies were especially important in the context of the population under study, since the extension of family allowances (cash transfers which were made only to formal-sector workers, i.e., those with higher incomes) to the poor was a controversial issue in Argentina at the time the survey was implemented in March 2009, a few months before a national midterm election. Opposition political parties, unions, academics and non-governmental organizations (including the Catholic Church) campaigned intensively for this policy measure, which was finally implemented by the federal government in November 2009 (Cruces and Gasparini, 2010). The public debate at the time of the survey focused on this specific policy measure and on the details of its coverage and implementation. Finally, apart from their salience in the political debate, another advantage of studying these policies is that government direct transfer programs can reach a substantial proportion of the population and can be fairly accurately targeted to the poor (unlike most of Argentina's public expenditure; see Gasparini and Cruces, 2010). This implies that low-income respondents faced a significant probability of being directly affected by policy changes in this area.[19]

The questionnaire was used to gather information on relevant socioeconomic characteristics of the respondents and their households and on their perceived and objective

[18] The information treatment in this context consists of a reduced number of sentences concerning the income distribution and a statement pointing out discrepancies. This is a substantially weaker treatment than that used in some other studies, such as the information provided about taxes and benefits in Chetty and Saez (2009). The significant results presented in the following pages are thus all the more remarkable, since they stem from a relatively limited treatment.

[19] The most recent experience was a well-targeted emergency program implemented in 2002, which covered 20 percent of the households in the country. The new family allowances program reached the same level of coverage and targeting in 2010, a year after its implementation (see Cruces and Gasparini, 2010, for a detailed account of both programs). The public debate on social policy at the time of the survey thus covered issues that could potentially have a direct impact on lower-income households.

positions in the distribution. The post-treatment questions[20] in the survey were designed to capture respondents' views on some of the prominent aspects of the public debate on distributional policy changes: whether to provide transfers to the poor or not and, if so, their modality (in cash, in kind or in some form of employment intermediation). The main variable for the analysis consisted of an indicator of strong pro-redistribution attitudes, which was equal to 1 for respondents who supported these three forms of government assistance simultaneously. Only 12.8 percent of the respondents in the overall sample (without distinguishing between the treatment and control groups) exhibited this high degree of support for redistribution. As expected, this proportion decreased monotonically with income levels, from a high of 21.2 percent for the poorest quintile of the sample to only 7 percent for the top quintile.

With respect to the components of this composite variable, the survey respondents were first asked: *"Do you think that the government should help poor people by giving them money?"* This question was answered positively by 14.7 percent of the interviewees (ranging from 23.2 in the bottom quintile to 9.5 percent in the top quintile). The survey also included what was essentially the same question but in reference to food. This type of in-kind transfer was supported by 33.5 percent of respondents, with the percentage also decreasing monotonically by income level (42.1 to 22.1 percent for the same quintiles as above). Finally, respondents were asked if the government should help the poor "*by providing them with a job*"; this option was supported by 98 percent of the respondents, with a virtually constant proportion across income groups.[21] The correlation between the answers to the questions about money and food is 0.44 and significant at the 1 percent level, which indicates that the joint analysis of these responses (in the composite variable) is warranted. In addition, they should also be studied independently of one another, since a substantial proportion of the information provided by the responses to each of these questions does not overlap.[22]

While the experiment consisted of providing information to the respondent, its setup provided for three alternative treatments. For individuals with positive biases (those who perceive themselves as being above their actual objective position), the interviewer

[20] Table A2 in the appendix presents the differences in pre-treatment variables (i.e., the questions asked before the intervention by the interviewer in the treatment questionnaire) by treatment status. A simultaneous test indicates that these differences are not significantly different from zero.

[21] While the low level of variation in this variable implies that it cannot be studied independently, it is included as a component of the overall "strong support for redistribution/government assistance to the poor" indicator.

[22] There are several reasons for the divergence between responses to the questions concerning the provision of assistance in the form of money and in the form of food. Most notably, respondents tend to prefer in-kind transfers (such as food transfers) because of paternalistic concerns (i.e., the belief that the poor might not spend cash on the "right" goods).

highlighted the fact that, according to the existing estimates, there were actually fewer households with a lower income level (i.e., they had overestimated their rank). On the other hand, individuals with negative biases (those who perceive themselves as being below their actual objective position) were informed that there were actually more households with a lower income than theirs (i.e., their position in the distribution was higher than they had thought). Finally, for those whose estimation coincided with the existing estimates, the interviewer simply pointed this out. These three treatment groups must be analyzed separately, because they provided the respondent with different types of information: overestimation, underestimation and confirmation of the respondent's own position.

In view of the discussion in Section 2, these treatments might be expected to have substantially different impacts on stated attitudes towards redistribution. Self-interest should induce those who overestimated their own rank to demand more redistribution (in the form of government assistance to the poor), since the treatment will make them more likely to consider themselves as potential beneficiaries of this type of policy. For those who underestimated their rank, the treatment implies that they will learn that they are richer and thus that they are less likely to benefit from the help towards the poor, so self-interest should induce them to demand less redistribution (although, as discussed later, altruistic concerns can override this effect). Finally, for respondents with no bias, the treatment simply confirms their perception of their position in the distribution; they do not receive any new information, and thus the treatment should be immaterial to their political attitudes. An impact of the treatment for this group would imply an effect of the interviewer's statement independent of its content, which could be difficult to separate from that of the actual information provided.[23]

Table 3 presents the mean levels of the preferences for redistribution variables (*strong support*, *help with money* and *help with food*) and a fourth dependent variable, which equals 1 if the respondent reports having made a donation in the recent past (this is used as a false experiment, as discussed below.) The results are presented for each of the treatment subgroups: those with a strictly negative bias, those with a zero bias, and those with a strictly positive bias. The first two lines in each panel display the means for each dependent variable by treatment status. The following line shows the difference between the two and the p-value of the mean difference test. Finally, the last line in each panel presents the results from the regression version of the test, i.e., the coefficient of the treatment indicator in an OLS

[23] An effect of the treatment for this group could also indicate that the respondent is misreporting his or her income level, in which case the true informational content of the interviewer's statement would not necessarily be a confirmation of the respondent's self-assessment. The fact that there are no significant effects for this group suggests that under-reporting is not the driving force behind the revealed biases.

regression of the dependent variable, which also includes the individual and household controls of columns 5 and 6 of Table 2, as well as neighborhood fixed effects and clustering of standard errors at the neighborhood level. More detailed information on these controls is provided in the note to the table.[24]

The results shown in column 2 of Table 3 indicate that, as predicted, the treatment had no statistically discernible effect among those with a zero bias regarding any of the variables. Column 1 of Table 3 presents the estimates for individuals with negative biases. These results also indicate that the treatment had no statistically significant effect on any of the outcomes that were analyzed. Finally, the results for individuals with positive biases (those who overestimate their rank) in column 3 point to a series of statistically significant effects. Treated individuals in this group exhibited a substantially higher proportion of strong support for redistribution (4.4 percentage points higher), support with money (3.5 percentage points) and support with food (5.1 percentage points). While the raw differences are just above standard significance levels, the conditional treatment effects are all positive and significant (at the 10 percent level for *strong support* and *help with money* and at the 5 percent level for *help with food*).

This effect of the information treatment on respondents who overestimated their income rank is both statistically and economically significant. It is also robust to a series of tests, as shown in columns 4, 5 and 6. A first robustness test involves using not only the sign of the bias, but also its intensity: the fifth column excludes those who misperceived their correct position by only one decile from those with a positive bias. The remaining individuals were informed that a substantially larger gap existed between their perceptions and reality, and the results should therefore be stronger. The evidence set forth in Table 3 indicates that this is indeed the case: the conditional and unconditional differences are substantially higher, and they are significant even for the unconditional t-tests. Columns 5 and 6 present two versions of a further robustness check. Individuals who consider their households to be poor should exhibit greater support for policies to help the poor in any case – this effect is present in the data for the three outcomes considered – and thus should be less affected by information that tells them that they are poorer than they thought. On the other hand, individuals who consider their households to be non-poor or relatively well-off should be more affected by the news that their relative position is lower than they expected. They may not want the government to help the poor so long as they do not consider themselves to be

[24] As in other small-sample experimental studies, the inclusion of control variables reduces the variability of the error term, which increases the statistical power of the significance test of the treatment effect.

part of this group, but an innovation in this sense could prompt them to favor redistribution out of self-interest, since they will realize that they have a greater probability of benefiting from this policy. Column 5 reports the treatment effect for the subset of individuals with positive biases who originally thought that their households had incomes equal to or above the median income for the general population, while column 6 conditions on those with no self-perception of poverty (which was captured in the pre-treatment phase of the survey). As expected, the effects are stronger and more significant for both groups than for those with a positive bias.[25]

Finally, the bottom rows in Table 3 present a false experiment designed to capture the presence of spurious effects of the treatment on respondents. These are estimates of treatment effects for the variable defined by the post-treatment question: *"Have you made any donations to an individual or charity during the past 12 months?"* The treatment may have an effect on stated preferences or demands, but cannot affect past actions. This donation variable was chosen because of its strong relationship to assistance to the poor and redistribution in general. If, instead of an increased willingness to redistribute, the provision of information only manages to induce more caring or generous statements from the respondent, then, by a shaming effect or by modifying the perceptions of the distribution, the treatment could also have a spurious impact on statements about what the respondent did in the past. The results shown in the last rows of Table 3 indicate that there is no significant effect of this type for any of the groups considered.

To sum up, as predicted, the treatment did not have an impact on individuals with no bias in their perceptions.[26] The lack of significant results for those with negative biases may reflect the offsetting effects on preferences from self-interest and altruistic mechanisms. As discussed in Section 2, individuals with a negative bias not only underestimate their own income rank, but they also underestimate the number of persons who are poor. As a consequence, the treatment informs them that there is more poverty than they originally thought, which can increase support for redistribution and then offset the selfish channel. Finally, respondents who were informed that they were relatively poorer than they had thought became more supportive of redistribution to the poor when informed about their true

[25] The results are qualitatively similar when conditioning the group with a positive bias to those who thought that their households belonged to the lower class or to those who reported having friends who were mainly from the lower class.

[26] Moreover, the results of the experiment also back up the findings about the relevance of reference groups for the formation of distributional perceptions that were presented in Section 4.2. If individuals know their true rank in the distribution but misreport it in the survey (because of embarrassingly low or high relative levels) or report focal-point answers because of a lack of interest, the provision of information should have no effect on their stated preferences, since they already have this unbiased estimate.

income rank. These results are sizeable: the difference in the support for redistribution between treatment and control group amounts roughly to half the difference in those variables between the top and bottom quintiles of the sample.

5 Conclusion

The motivation of this paper was the lack of evidence on the accuracy and origins of perceptions of the income distribution, which play a crucial (though implicit) role in political economy and public finance models. The study proposed a simple mechanism where agents extrapolate from their reference group without accounting for the selection process that led to their group – either because of informational or cognitive limitations. A tailored household survey provided evidence on the salient role of reference groups for the formation of subjective distributions, which are significantly correlated with the respondent's relative position within her reference group (as proxied by area of residence), and on the presence of sizeable systematic biases in these perceptions in Argentina. The systematic biases documented in this paper and their consequences, however, can arise in any society – the determinant factor is the sorting of individuals in reference groups by income level. For instance, the posited mechanism for the formation of biased perceptions is compatible with the systematic downward bias in the perception of wealth inequality in the United States documented by Norton and Ariely (2011).

Furthermore, this study implemented an original survey experiment in the field, where a randomly assigned group of respondents were provided consistent information about their ranking in the income distribution as a feedback to their responses. Confronting agents' biased perceptions with this information had a significant effect on stated preferences for redistribution: those who overestimated their relative position (who thought of themselves to be relatively richer than they were) demanded more redistribution. To the degree that the information treatment managed to correct biased distributional perceptions, these results can be interpreted as evidence of the causal effect of misperceptions on political attitudes. This mechanism provides an alternative explanation for the low levels of redistribution observed in modern democracies, and since it affects relatively poorer individuals it is reminiscent of the Marxian notion of false consciousness.[27] Having accurate information about the income distribution might induce agents to better calibrate their demands for redistribution. The

[27] Olin Wright's (2009) discussion of false consciousness states: "Ideology is seen as preventing workers from understanding the nature of their oppression and the possibilities of its transformation. The absence of effective struggle for socialism, then, is at least in part explained by the pervasiveness of these cognitive distortions."

results in this paper support Romer's (2003) discussion of the possible welfare improving effects of subsidizing information, and Besley's (2007) remarks about the potential of information providers for improving policies, although the impact of the biases in the efficiency of redistribution should also be considered (Acemoglu and Robinson, 2005),

The role of misconceptions in political economy has been studied before (Romer, 2003; Slemrod, 2006). While Besley (2007) highlights the benefits from incorporating notions of dispersed and limited information for modern political economy, building-in more specific factors like biased perceptions of the distribution can further enrich political economy models and empirical applications. It can also provide explanations for other puzzles in the literature, such as Bartels' (2005) findings on the support of low income voters for tax cuts to the rich and the dissonance between their views on inequality and public policy, or Bartels' (2008) results on the reduced responsiveness of representatives to low income voters. More generally, concepts such as inequality, self-interest and the median voter can be adapted in their application to political economy outcomes when misperceptions and misconceptions play a role.

The findings in this paper indicate effects of perceptions on stated preferences for redistribution. Further research could focus on the impact of biases and information on actual behavior – for instance, on charitable donations or on voting patterns. Moreover, the results in this paper could originate in either limited information or limited cognitive ability – further research could disentangle the source of the observed biases in distributional (and other) perceptions.

References

Acemoglu, D. and J. Robinson (2001), "Inefficient Redistribution," *American Political Science Review*, Vol. 95 (3), pp. 649-661.

Alesina, A. and E. Glaeser (2004), *Fighting Poverty in the US and Europe. A World of Difference*, Oxford University Press.

Alesina, A. and G.M. Angeletos (2005), "Fairness and Redistribution: US vs. Europe," *The American Economic Review*, vol. 95, pp. 913-35.

Alesina, A. and P. Giuliano (2009), "Preferences for Redistribution", NBER Working Paper No. 14825.

Alesina, A. and E. La Ferrara (2005), "Preferences for Redistribution in the Land of Opportunities", *Journal of Public Economics*, 89(1) pp. 897-931.

Amiel, Y. and F. Cowel (1992), "Measurement of Income Inequality: Experimental Test by Questionnaire", *Journal of Public Economics*, vol. 47, pp. 3–26.

Bartels, L. (2005), "Homer gets a tax cut: Inequality and public policy in the American mind," *Perspectives on Politics*, Vol. 3 (1).

Bartels, L. (2008), *Unequal Democracy: The Political Economy of the New Gilded Age*, Russell Sage Foundation and Princeton University Press.

Bayer, P., S. Ross and G. Topa (2008), "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes", *Journal of Political Economy*, vol. 116 (6), pp. 1150-1196.

Bénabou, R. and E. Ok (2001), "Social Mobility and the Demand for Redistribution: The POUM Hypothesis", *Quarterly Journal of Economics*, vol. 116 (2), pp. 447-487.

Bénabou, R. and J. Tirole (2006), "Belief in a Just World and Redistributive Politics", *Quarterly Journal of Economics*, vol. 121(2), pp. 699-746.

Benoît , J.P. and J. Dubra (2011) "Apparent Overconfidence," *Econometrica*, forthcoming.

Bertrand, M. and S. Mullainathan (2001), "Do People Mean What They Say? Implications for Subjective Survey Data", *American Economic Review*, vol. 91(2), pp. 67-72.

Besley, T. (2007), "The New Political Economy," *Economic Journal*, vol. 117 (524).

Card, D., A. Mas, E. Moretti and E. Saez (2010), "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction", NBER Working Paper No. 16396.

Camerer, C., G. Loewenstein and M. Rabin (2003), *Advances in Behavioral Economics*, Princeton University Press.

Chetty, R. and E. Saez (2009), "Teaching the Tax Code: Earnings Responses to an Experiment with EITC Claimants", NBER Working Paper No. 14836.

Clark, A., P. Frijters and M. Shields (2008), "Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles", *Journal of Economic Literature*, vol. 46(1), pp. 95-144.

Cowell, F. (2000), "Measurement of Inequality", in A. Atkinson and F. Bourguignon (eds.) (2002), *Handbook of Income Distribution*, edition 1, vol. 1, Elsevier Press.

Cowell, F. and G. Cruces (2004), "Perceptions of Inequality and Risk", *Research on Economic Inequality*, vol. 12, pp. 99-132.

Cruces, G. and L. Gasparini (2010), "Las Asignaciones Universales por Hijo en Argentina: Impacto, Discusión y Alternativas", *Económica*, vol. 56 (1), pp. 105-146.

DellaVigna, S. (2009), "Psychology and Economics: Evidence from the Field", *Journal of Economic Literature*, vol. 47, pp. 315-372.

Di Tella, R., S. Galiani and E. Schargrodsky (2008), "Reality versus Propaganda in the Formation of Beliefs about Privatization", NBER Working Paper No. 14483.

Duflo, E. and E. Saez (2003), "The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment", *Quarterly Journal of Economics*, vol. 118 (3), pp. 815-842.

Fong, C. (2001), "Social Preferences, Self Interest and the Demand for Redistribution", *Journal of Public Economics*, vol. 82 (2), pp. 225-246.

Gasparini, L., W. Sosa Escudero, M. Marchionni and S. Olivieri (2008), "Income, Deprivation, and Perceptions in Latin America and the Caribbean: New Evidence from the Gallup World Poll", unpublished working paper, CEDLAS-UNLP and IDB.

Gasparini, L. and G. Cruces (2010), "A Distribution in Motion: The Case of Argentina", in Luis Felipe López Calva and Nora Lustig (eds.), *Declining Inequality in Latin America: A Decade of Progress?* Brookings Institution Press.

Georgiadis, A. and A. Manning (2011), "Spend It Like Beckham? Inequality and Redistribution in the UK, 1983-2004," *Public Choice*, forthcoming.

Glaeser, E., M. Resseger and K. Tobio (2008), "Urban Inequality", NBER Working Paper No. 14419.

Grether, D. M. (1980), "Bayes' Rule as a Descriptive Model: The Representative Heuristic", *Quarterly Journal of Economics*, vol. 95, pp. 537-557.

Grether, D. M. (1990), "Testing Bayes Rule and the Representativeness Heuristic: Some Experimental Evidence", *Journal of Economic Behavior and Organization*, vol. 17, pp. 31-57.

Hainmueller, J. and M. Hiscox (2010), "Attitudes toward Highly Skilled and Low-Skilled Immigration: Evidence from a Survey Experiment", *American Political Science Review*, vol. 104, pp. 61-84.

Horiuchi, Y., K. Imai and N. Taniguchi (2007), "Designing and Analyzing Randomized Experiments: Application to a Japanese Election Survey Experiment", *American Journal of Political Science*, vol. 51, No. 3, pp. 669-687.

Hurd, M. (2009), "Subjective Probabilities in Household Surveys", *Annual Review of Economics*, vol. 1, pp. 543-564.

Jensen, R. (2010), "The Perceived Returns to Education and the Demand for Schooling", *Quarterly Journal of Economics*, vol. 125 (2), pp. 515-548.

Joliffe, D. (2001), "Measuring Absolute and Relative Poverty: The Sensitivity of Estimated Household Consumption to Survey Design", *Journal of Economic and Social Measurement*, vol. 27(1/2), pp. 1-23.

Kahneman, D. and A. Tversky (1972), "Subjective Probability: A Judgment of Representativeness", *Cognitive Psychology*, No. 3, pp. 430-451.

Kahneman, D., P. Slovic and A. Tversky (eds.) (1982), *Judgment Under Uncertainty: Heuristics and Biases*, Cambridge University Press.

Luttmer, E. (2001), “Group Loyalty and the Taste for Redistribution”, *Journal of Political Economy*, vol. 109 (3), pp. 500-528.

Manski, C. (2004), “Measuring Expectations," *Econometrica*, vol. 72 (5), pp. 1329-1376.

Meltzer, A. and S. Richard (1981), “A Rational Theory of the Size of Government,” *Journal of Political Economy*, 89(5), pp. 914–927.

Norton, M. and D. Ariely (2011), “Building a Better America—One Wealth Quintile at a Time”, *Perspectives on Psychological Science*, vol. 6, pp. 9-12.

Nuñez, J. (2005), “Living Under a Veil of Ignorance”, mimeo, Universidad de Chile.

Olin Wright, E. (2009), “Class, State and ideology: An introduction to social science in the Marxist tradition”, Sociology 621 unpublished lecture notes, Department of Sociology, University of Wisconsin, Madison.

Olken, B. (2009), “Corruption Perceptions vs. Corruption Reality”, *Journal of Public Economics*, vol. 93 (7-8), pp. 950-964.

Piketty, T. (1995), “Social Mobility and Redistributive Politics”, *Quarterly Journal of Economics*, vol. 110 (3), pp. 551-584.

Rabin, M. (1998), “Psychology and Economics”, *Journal of Economic Literature*, vol. 36 (1), pp. 11-46.

Ravallion, M. and M. Lokshin (2002), “Self-Rated Economic Welfare in Russia”, *European Economic Review*, vol. 46 (8), pp. 1453-1473.

Romer, P. (2003), “Misconceptions and Political Outcomes”, *Economic Journal*, vol. 113, No. 484, pp. 1-20.

Rotemberg, J. (2002), “Perceptions of Equity and the Distribution of Income”, *Journal of Labor Economics*, vol. 20 (2), part 1.

Simon, H. (1972), *Theories of Bounded Rationality*, New York.

Slemrod, J. (2006), “The Role of Misconceptions in Support for Regressive Tax Reform”, *National Tax Journal*, vol. 59, pp. 57–75.

Topa, G. (2001), “Social Interactions, Local Spillovers and Unemployment”, *Review of Economic Studies*, vol. 68 (2), pp. 261-295.

Tversky, A. and D. Kahneman (1974), “Judgment under Uncertainty: Heuristics and Biases”, *Science*, vol. 185, pp. 1124-1131.

Zafar, B. (2011), “How Do College Students Form Expectations?,” *Journal of Labor Economics*, vol. 29 (2), pp. 301-348.

Tables

Source for all tables: Own calculations based on the *Survey on Distributional Perceptions and Redistribution*.

Table 1. Objective income decile, perceived own income decile and bias by quintile of objective income

Quintiles of population income	Average objective decile	Average perceived own decile	Mean bias	Proportion with positive bias	Average positive bias	Proportion with negative bias	Average negative bias
Lowest	1.62	4.60	2.98	0.85	3.02	0.04	-0.04
Second	3.47	4.96	1.49	0.71	1.71	0.16	-0.21
Third	5.53	5.38	-0.14	0.30	0.60	0.40	-0.74
Fourth	7.54	5.89	-1.64	0.07	0.09	0.81	-1.73
Highest	9.35	6.48	-2.88	0.00	0.00	0.97	-2.88
Total (N=1060)	6.12	5.60	-0.53	0.30	0.75	0.55	-1.28

Notes: the bias is defined as the perception of income decile minus objective income decile (see Table A1 for detailed definitions).

Table 2. Determinants of perceived own income decile

	Dependent variable: Perceived own income decile					
	(1)	(2)	(3)	(4)	(5)	(6)
Objective income decile	0.2452	0.2099	-0.0168	F-test†	0.0048	0.0109
	[0.0245]***	[0.0280]***	[0.0944]		[0.1237]	[0.1228]
Rank within locality	–	–	0.2151	0.2311	0.2002	0.2288
			[0.0868]**	[0.1195]*	[0.1114]*	[0.1096]**
Has friends from all social classes	–	–	–	–	–	0.5046
						[0.2897]*
Interaction: Locality rank & friends variable	–	–	–	–	–	-0.0937
						[0.0458]**
Constant	4.0916	3.8997	4.2961	4.2846	4.1199	3.976
	[0.1798]***	[0.2266]***	[0.2659]***	[0.4379]***	[0.6554]***	[0.6709]***
Observations	1054	1054	1054	1054	1045	1045
R-squared	0.12	0.18	0.19	0.19	0.22	0.22
Neighborhood fixed effects	No	Yes	Yes	Yes	Yes	Yes
Individual controls	No	No	No	No	Yes	Yes
Levels of objective decile as indicators	No	No	No	Yes	No	No

Notes: Standard errors clustered by neighborhood in brackets. * significant at 10%; ** significant at 5%; *** significant at 1%. † F-test of joint significance for nine objective income decile indicator variables: p-value of 0.2468. The individual controls in the regressions in columns 5 and 6 include the sex of the respondent, whether the respondent is the household head, his or her age, indicators for his or her education level (from primary incomplete and lower up to postgraduate degree), whether the respondent has a spouse, indicators for the spouse's education level (if present), whether the respondent is a public employee, whether the respondent is unionized, and whether the household has any extra source of income besides labor earnings.

Table 3. Biased perceptions of income distribution and preferences for redistribution: Experimental results.

	Negative bias: Treatment tells respondents that their position is higher than they had estimated	No bias: Treatment confirms respondents' positional perception	Positive bias: Treatment tells respondents that their position is lower than they had estimated			
			All	Bias>1	Self-perception: at or above median income	Self-perception: not poor
	(1)	(2)	(3)	(4)	(5)	(6)
	Government should provide assistance to the poor (strong support)					
Treatment group	0.100	0.186	0.194	0.200	0.212	0.240
Control group	0.087	0.153	0.148	0.122	0.129	0.161
Difference [p-value]	0.013 [0.708]	0.033 [0.58]	0.045 [0.145]	**0.078 [0.058]**	**0.084 [0.033]**	**0.079 [0.093]**
Conditional diff. [p-value]	0.026 [0.331]	0.025 [0.813]	**0.074 [0.093]**	**0.113 [0.067]**	**0.112 [0.030]**	**0.145 [0.024]**
	Government should help the poor with money					
Treatment group	0.111	0.226	0.212	0.242	0.227	0.250
Control group	0.108	0.153	0.176	0.150	0.152	0.153
Difference [p-value]	0.003 [0.553]	0.073 [0.246]	0.035 [0.217]	**0.092 [0.045]**	**0.074 [0.060]**	**0.097 [0.054]**
Conditional diff. [p-value]	0.010 [0.637]	0.063 [0.601]	0.084 **[0.075]**	**0.132 [0.056]**	**0.131 [0.013]**	**0.186 [0.005]**
	Government should help the poor with food					
Treatment group	0.284	0.381	0.424	0.444	0.422	0.413
Control group	0.303	0.347	0.373	0.381	0.362	0.376
Difference [p-value]	-0.019 [0.304]	0.034 [0.663]	0.051 [0.180]	0.064 [0.173]	0.060 [0.160]	0.037 [0.310]
Conditional diff. [p-value]	-0.024 [0.484]	-0.010 [0.928]	**0.118 [0.014]**	**0.124 [0.092]**	**0.149 [0.032]**	**0.132 [0.099]**
	False experiment: Made donations in the last twelve months					
Treatment group	0.866	0.788	0.719	0.687	0.715	0.755
Control group	0.831	0.845	0.742	0.722	0.764	0.793
Difference [p-value]	0.035 [0.883]	-0.057 [0.364]	-0.023 [0.675]	-0.035 [0.712]	-0.049 [0.820]	-0.038 [0.728]
Conditional diff. [p-value]	0.044 [0.155]	-0.024 [0.777]	0.029 [0.593]	0.031 [0.632]	-0.004 [0.951]	0.042 [0.604]
N:	573-586	156-158	304-308	212-215	266-272	177-183

Notes: p-value of differences in brackets; numbers in bold indicate significance at the 10 percent level or lower. The test for the unconditional difference reports the p-value for $\mu_T<\mu_C$ for cases of negative bias (column 1), $\mu_T=\mu_C$ for those with no bias (column 2) and for $\mu_T>\mu_C$ for those with positive bias (columns 3-6). The conditional difference is computed from a regression of the outcome of interest against a treatment indicator, neighborhood fixed effects and a series of individual controls. The conditional difference is the estimate of the coefficient of the treatment indicator, and the reported p-value (derived from standard errors clustered at the neighborhood level) is the significance of this coefficient. The individual controls in the regressions include the sex of the respondent, whether the respondent is the household head, his or her age, indicators for his or her education level (from primary incomplete and lower up to postgraduate degree), whether the respondent has a spouse, indicators for the spouse's education level (if present), whether the respondent is a public employee, whether the respondent is unionized, and whether the household has any extra source of income besides labor earnings. The bias is defined as the perception of income decile minus objective income decile. "Self-perception: median or above" means that the perceived income decile is 6 or higher. "Self-perception: not poor" means that the individual did to consider his or her household to be poor (this question preceded the treatment in the questionnaire). See Table A1 for further variable definitions.

Figures

Source (except Figure 1): Own calculations based on the *Survey on Distributional Perceptions and Redistribution*.

Figure 1. Illustration of differences between population and reference group income distributions

1a: Biases with a rich reference group 1b: Biases with a middle-class reference group

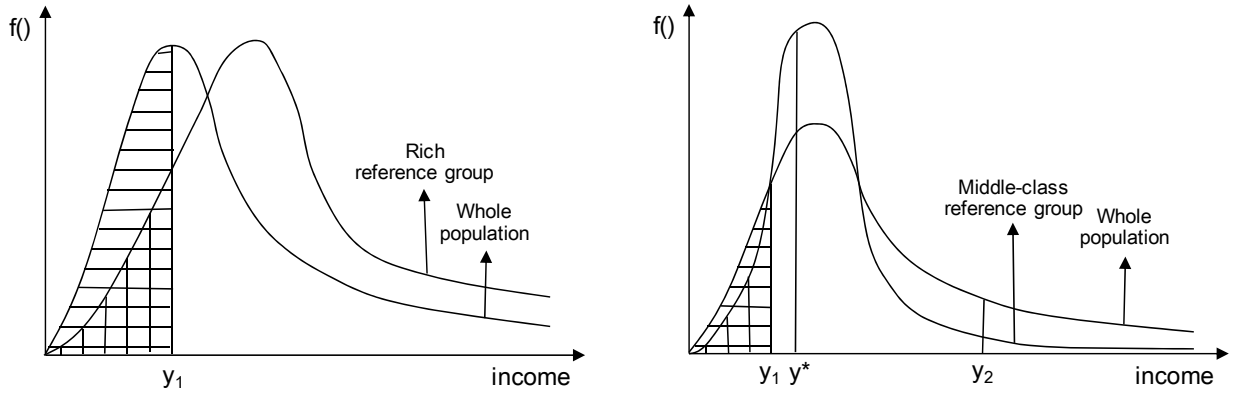

Figure 2. Distribution of objective and perceived own income decile

2a 2b

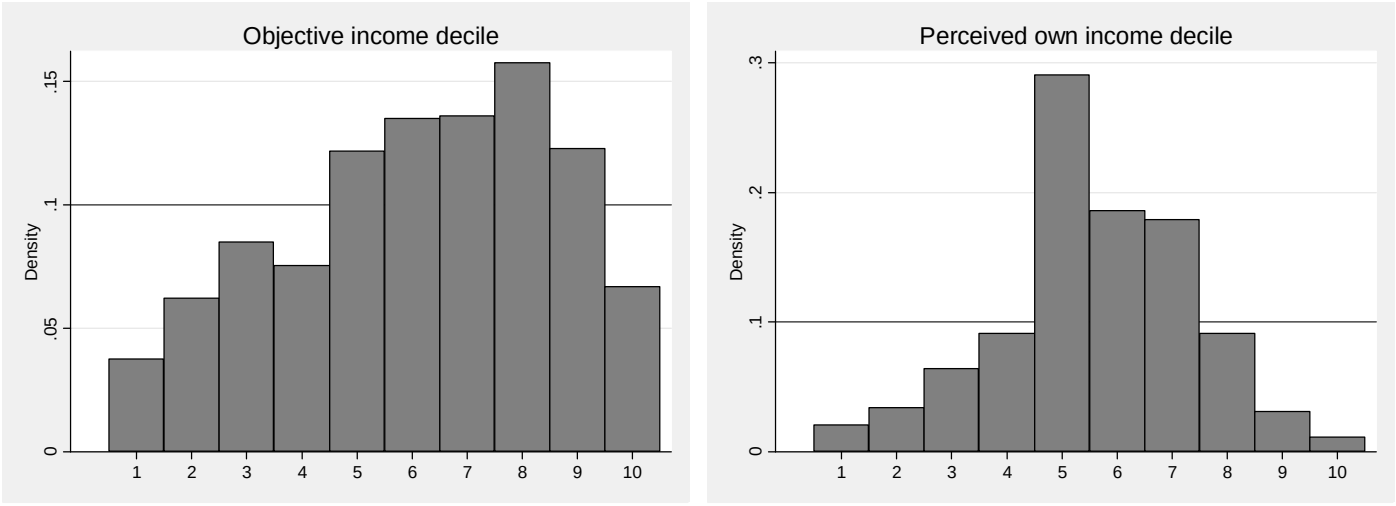

Notes: See Table A1 for further definitions. Obs: 1,060.

Figure 3. Distribution of bias in perception of own income decile, in number of deciles

3a 3b

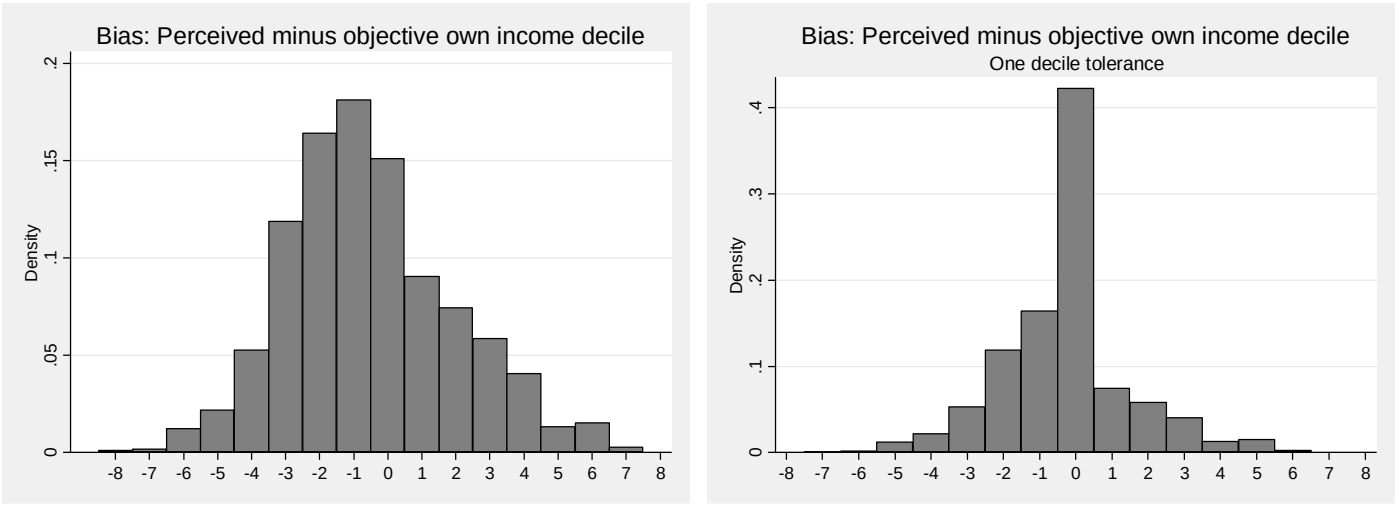

Notes: bias is defined as the perception of income decile minus objective income decile. See Table A1 for further definitions. Obs: 1,060.

Figure 4. Average perception of own income decile and bias, by objective income decile

4a 4b

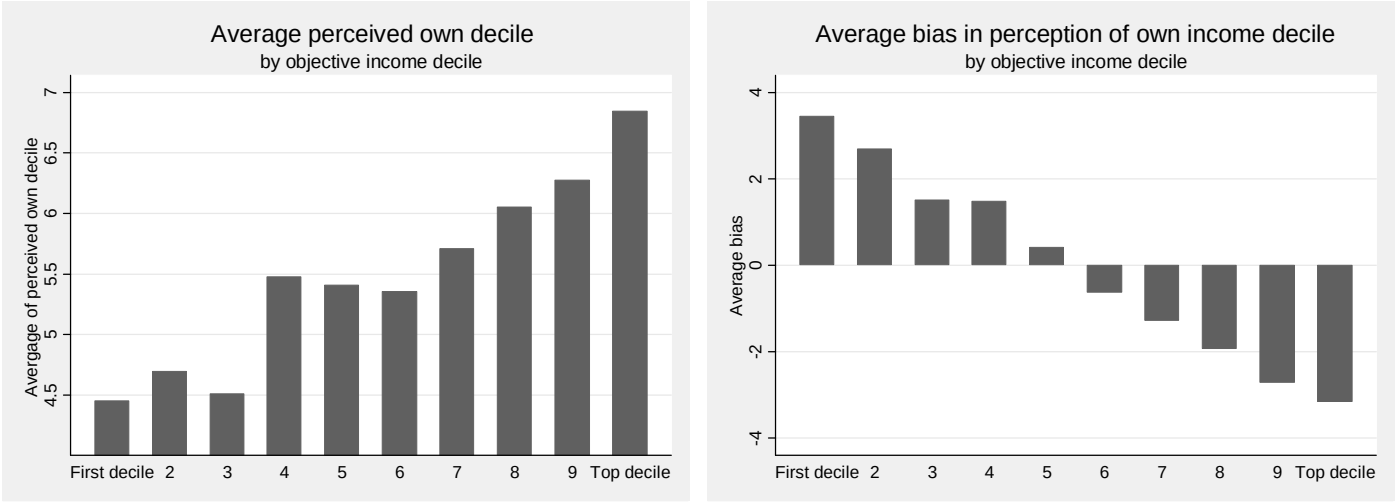

Notes: bias is defined as the perception of income decile minus objective income decile. See Table A1 for further definitions. Obs: 1,060.

Figure 5. Average objective decile, perceived own decile and bias by neighborhood

5a 5b

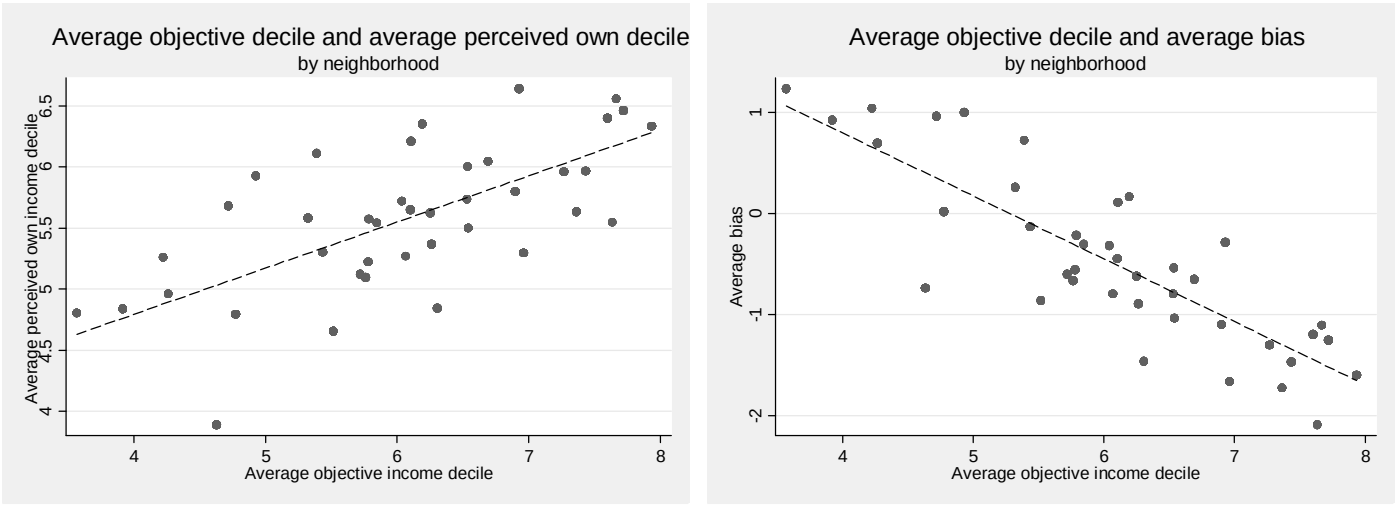

Notes: The average is taken over 41 neighborhoods. Bias is defined as the perception of income decile minus objective income decile. See Table A1 for further definitions. Obs.: 1,060.

Appendix tables: Table A1. Variable definitions and descriptive statistics

Variable	Description - relevant question from the *Survey on Distributional Perceptions and Redistribution* March 2009, Greater Buenos Aires, Argentina	Mean	SD	Min.	Max.	Obs.
Objective income decile	The interviewer displayed a table with income ranges computed by the researchers, corresponding to the deciles of the distribution of total household income for Argentina at the time of the survey. Question: *I will show you a Table with levels of income. Please indicate where, approximately, would you say that your family is located, considering all income in your household from every concept (work, government transfer programs, pensions, rent, etc.) (1) Less than X; (2) X to Y; ... ; (10) More than Z.*	6.122	2.462	1	10	1060
Perception of own income decile	The interviewer made a statement and asked the following question: *There are 10 million families in Argentina. Of those 10 million, how many do you think have an income lower than yours? (1) 0 to 1 million; (2) between 1 and 2 millions; ... ; (10) 9 to 10 millions.*	5.595	1.773	1	10	1060
Bias	The bias is constructed as the level of the *objective income decile* minus that of the *perceived own income decile*. It is negative for those who consider themselves to be in a lower position than they really are, and it is positive for those who consider themselves to be in a higher position than they really are.	-0.526	2.489	-8	7	1060
Treatment	Half of the sample was assigned to a "treatment" questionnaire with the following specific intervention from the interviewer, which was not present in the "control" version. The interviewers alternated questionnaire types. The intervention consisted of comparing the answer from the *objective income decile (X)* to that of the *perception of own income decile (Y)*, and stating accordingly: The interviewer read the following statement (with X and Y being determined by previous answers): *"Based on your income level, the latest studies conducted by the University indicate that there are X million families with an income lower than yours, while you stated that there were Y."* The interviewer then read out one of the three following statements, depending on the accuracy of the X/Y comparison: (1) *"In fact, there are more families with a lower income than yours than you believed"*, (2) *"You were right about how many families have a lower income than yours"*, or (3) *"In fact, there are fewer families with a lower income than yours than you believed".*	0.514	0.500	0	1	1060
Rank within locality	This variable is constructed using the *objective income decile* variable for each respondent and computing her ranking within her area of residence, where the 41 sampling points where aggregated to 10 geographic localities. The rank within the locality is computed using all observations in each area in the sample, and transformed to a 1-10 scale (as the objective and subjective income decile variables). It is computed as the number of households with lower income than that of the respondent divided by the total number of households.	5.419	2.483	1.083	9.779	1060
Respondent has friends from all social classes	This is an indicator variable equal to one if the respondent states that he or she has friends from all social classes when asked the question: *"Among your friends and co-workers, would you say that there are individuals from all social classes (1), or, if not, that most of your friends belong to the lower class (2), the middle class (3), or the upper class (4)?"*	0.376	0.485	0	1	1060
Self-perception of poverty	This is an indicator variable equal to one if the respondent answered "Yes" to the question: *"Do you think that your family is poor? (1) Yes; (0) No."*	0.224	0.417	0	1	1049
Help the poor with money	*"Do you think that the government should help poor people by giving them money? (1) Yes; (0) No."*	0.148	0.355	0	1	1049
Help the poor with food	*"Do you think that the government should help poor people by giving them food? (1) Yes; (0) No."*	0.336	0.472	0	1	1049
Help the poor find jobs	*"Do you think that the government should help poor people by helping them to find a good job? (1) Yes; (0) No."*	0.982	0.133	0	1	1052
Government should provide assistance to the poor	Indicator variable constructed as a function of the three previous questions. Equal to 1 if the respondent answered affirmatively to these questions, and equal to 0 otherwise.	0.129	0.336	0	1	1060
Respondent made donations in the past 12	*"Have you donated money, food or clothes to any charity or individual in need in the past twelve months? Yes (1); No (0)"*	0.808	0.394	0	1	1045
Sex	(1) Male; (0) Female.	0.489	0.500	0	1	1060
Age	Age in years.	49.2	15.4	17	88	1051
Educational level of the respondent or his/her spouse	Indicator variables for the following categories: *(1) Primary incomplete; (2) Primary complete; (3) Secondary incomplete; (4) Secondary complete; (5) Undergraduate incomplete; (6) Undergraduate complete; (7) Postgraduate.* This table reports the average of these categories for the respondent.	4.049	1.599	1	7	1054
Household head	*"Are you the head of the household? (1) Yes; (0) No."*	0.757	0.429	0	1	1060

Note: the sample is restricted to the 1,060 observations with non-missing bias information (objective and perceived income decile), which corresponds to the sample analyzed in the paper. Further notes below.

Notes to Table A1

The use of income intervals significantly reduces non-response rates, as shown in a 2007 pilot and in large-scale international projects such as the Gallup World Poll, which concentrate on total household income rather than on its components (Gasparini et al., 2008). The boundaries of the intervals corresponded to actual deciles of the distribution, which facilitated the comparison of objective and perceived rank implemented in the experimental design.

To ensure comparability between the objective and subjective income ranks, interviewers were instructed to impute the lowest category for respondents who considered that less than 1 million households had a lower income than theirs, the next-highest category for those who responded with any number between 1 and 2 million, and so forth until reaching the highest category (10) for those who reported any number between 9 and 10 million.

It should be noted that as a result of the Argentine government's intervention in the operations of the National Statistics and Census Institute (INDEC) in 2007, the availability of reliable household survey microdata and of official income distribution indicators was quite limited until 2010. To construct the deciles for 2009, the team updated the boundaries of total household monthly income deciles from 2007 using information from INDEC's monthly index of wage levels, which continued to be published. When the results of the 2009 national household survey became available in 2010, all of the estimated decile boundaries fell within the 95 percent interval of the actual points in the microdata.

The two versions of the questionnaire (and their English translations) are available upon request.

Table A2. Differences in pre-treatment variables between treatment and control groups

Variable	Treatment group	Control group	Difference (μ_T-μ_C)	t ratio of difference
Age	49.99	48.50	1.49	-1.6
Head of household indicator	0.760	0.755	0.005	-0.21
Male indicator	0.470	0.490	-0.020	-0.65
Number of adults living in the household	1.809	1.830	-0.021	-0.28
Number of children (14 and below) in household	0.737	0.739	-0.002	-0.04
Number of own children	2.032	1.917	0.115	-1.33
No spouse in household	0.112	0.133	-0.021	-1.06
Household has fixed phone line	0.804	0.832	-0.028	-1.22
Household rents dwelling	0.245	0.221	0.024	-0.94
Number of members working	1.571	1.569	0.002	-0.03
Household receives government transfers (welfare)	0.047	0.042	0.005	-0.39
Household has income sources besides labor earnings	0.079	0.070	0.009	-0.55
Some primary education (complete or incomplete)	0.228	0.216	0.012	-0.47
Some secondary education (complete or incomplete)	0.423	0.413	0.010	-0.34
Some higher education (complete, incomplete or postgraduate)	0.349	0.371	-0.022	-0.75
Housewife	0.169	0.152	0.017	-0.75
Wage earner	0.275	0.294	-0.019	-0.7
Liberal profession	0.166	0.149	0.017	-0.77
Pensioner	0.159	0.123	0.036	-1.71
Looking for a job	0.043	0.077	-0.034	-2.37
Working	0.654	0.689	-0.035	-1.24
Unionized	0.224	0.232	-0.008	-0.33
Public sector worker	0.140	0.138	0.002	-0.08
In formal employment	0.355	0.357	-0.002	-0.08
Perceives household as poor	0.198	0.238	-0.040	-1.62
Log of assessed minimum living income	8.128	8.108	0.020	-0.73
Has friends from all social classes	0.361	0.389	-0.028	-0.98
Perceived own income decile	5.676	5.514	0.162	-1.51
Objective income decile	6.193	6.021	0.172	-1.14

Notes: the table includes all 1,115 observations in the database, including those with incomplete or missing answers. Scheffe's method for simultaneous testing provides the critical t-statistic for the significance of each of the tests in the table. For a 95 percent level of significance, with tests and 1114 degrees of freedom, the critical value is 6.64. None of the differences is significantly different from zero according to this method.