



Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment[☆]

Guillermo Cruces^{a,b,c,*}, Ricardo Perez-Truglia^d, Martin Tetaz^a

^a Centro de Estudios Distributivos Laborales y Sociales [Center for Social, Labor and Distributional Studies] (CEDLAS), Universidad Nacional de La Plata, La Plata, Argentina

^b IZA: Institute for the Study of Labor, Bonn, Germany

^c CONICET [National Scientific and Technical Research Council], Argentina

^d Economics Department, Harvard University, MA, USA

ARTICLE INFO

Article history:

Received 17 May 2011

Received in revised form 8 October 2012

Accepted 23 October 2012

Available online 12 November 2012

JEL classification:

D31

D83

H24

H53

I30

Keywords:

Income distribution

Biased perceptions

Limited information

Preferences for redistribution

ABSTRACT

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 explores their potential impact on preferences for redistribution. A tailored household survey provides original evidence on systematic biases in individuals' evaluations of their own relative position in the income distribution. The study discusses one of the mechanisms that may generate such biases, based on the extrapolation of information from endogenous reference groups, and presents some suggestive evidence that this mechanism has significant explanatory power. The impact of these biased perceptions on attitudes toward redistributive policies is studied by means of an experimental design that was incorporated into the survey, which provided consistent information on the own-ranking within the income distribution to a randomly selected group of respondents. The evidence suggests that those who had overestimated their relative position and thought that they were relatively richer than they were tend to demand higher levels of redistribution when informed of their true ranking.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

The shape of the income distribution plays a key role in the determination of policies with redistributive components (such as those dealing with social security, health care, government transfers and taxation) in political economy models. However, the main policy determinant is not its actual shape, but rather how it is perceived by agents in the economy. Additionally, individuals' perceptions of the

income distribution can affect how they will react to redistributive policies (for instance, through tax evasion), which is a key input for public finance models. This study fills a gap in the literature by exploring the origins 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

[☆] 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 editor, Dan Silverman, and two anonymous referees provided thorough and constructive comments that were taken into account in this version of the paper. 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 about the project which led 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 Nopo, Michael Norton, Laura Ripani, Martín Rossi, László Sándor, 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. Funding for the *Encuesta de Percepciones Distributivas y Redistribución* was generously provided by 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.

* Corresponding author at: CEDLAS, Facultad de Ciencias Económicas, Universidad Nacional de La Plata, Calle 6 entre 47 y 48, 5to. piso, oficina 516 (1900), La Plata, Argentina. Tel./fax: +54 221 4229383.

E-mail addresses: gcruces@cedlas.org, cedlas@depeco.econo.unlp.edu.ar (G. Cruces).

URL: <http://www.cedlas.org> (G. Cruces).

contribution, see Piketty, 1995). For instance, when forming their views on public policies, agents may need to infer the importance of effort and of predetermined factors in the income generation process (Piketty, 1995; Bénabou and Tirole, 2006); in so doing, they may evaluate the prospects for 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 complicated 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 to explain how they are formed (Zafar, 2011). In an application to distributional issues, Norton and Arieli (2011) elicit information on Americans' perceptions of the wealth distribution in their society and find significant discrepancies between actual and perceived levels of inequality. This paper documents systematic differences between objective and subjective income distributions and sheds light on the origins of these discrepancies. Most importantly, an experimental design makes it possible to test whether correcting these biases has an impact on stated preferences for redistribution.

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 then infer the entire distribution from that information. If agents do not fully account for the selection process involved in the formation of the sample that they observe, their inferences will be systematically biased. This failure may be due to limitations in the available information set which arise from the fact that information may be costly or difficult to obtain. Alternatively, agents may have the necessary information, but they may sometimes fail to use it correctly, as argued in the cognitive bias literature. 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 same data can be used to study biases and preferences for redistribution, which is the main focus of this study.

The empirical results presented in this paper are based on the *Survey on Distributional Perceptions and Redistribution*, a study of 1100 representative households in Greater Buenos Aires 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 that systematic biases are present in perceptions of own income rank: a significant portion of poorer individuals place themselves in higher positions than they actually occupy, while a significant proportion of richer individuals underestimate their rank. Moreover, as predicted, the bias is significantly correlated with the respondents' relative positions within the reference group (as proxied by area of residence). Also in keeping with the posited mechanisms, respondents with friends from heterogeneous social backgrounds are less prone to these biases.

Finally, the study explores how these misperceptions about the income distribution may affect attitudes toward redistribution. For instance, self-interest might induce poor individuals to demand less redistribution if they think they are relatively richer than they actually are. This study 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). In addition, in this survey experiment, biased subjects were provided with feedback and were actually confronted with accurate information.

The results from the experiment indicate that confronting agents' biased perceptions with accurate information had a significant effect on their stated preferences for redistribution. Those who underestimated their income ranking did not change their attitudes toward redistribution when provided with accurate information about their income ranking. However, those who overestimated their relative position (i.e., who thought that they were relatively richer than they are) and who were provided with accurate 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. 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 then goes on to explore these factors' implications for attitudes toward redistribution. The third section describes the household survey and outlines the randomized experiment that was designed to answer these questions. The fourth section presents the empirical results on biased perceptions of income distribution, and the fifth section describes the identification strategy and the results from the experiment on biases and preferences for redistribution. The last section concludes.

2. Subjective income distributions, potential biases and preferences for redistribution

Economic agents' assessments of income distributions depend on their access to information and on their ability to process the relevant data. The latter is a trivial consideration 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. Agents can be deemed sophisticated if they apply Bayes' rule to infer the income distribution for the entire population from the subset that they observe. A naïve agent is denoted by a failure to fully apply Bayes' rule. This failure can result in biased perceptions of the overall income distribution.¹

An agent may arrive at naïve estimates under certain circumstances. First, the information about the income distribution may be costly to acquire, or the advantages of doing so may not be evident. It may 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. 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 be very imprecise. Indeed, the use of heuristics in statistical inference and the systematic biases that such an exercise entails is a well-documented phenomenon in the cognitive literature (Kahneman et al., 1982). The most relevant case in this discussion is the representativeness heuristic, in which individuals fail to apply Bayes' rule to the information they

¹ Cruces et al. (2011) provide a lengthier and more detailed discussion of the factors at work in the context of a statistical inference problem.

obtain (Kahneman and Tversky, 1972). 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 an extreme situation in which a naïve agent uses the information about the income distribution within his or her reference group as if it were representative of the entire population. If the formation of reference groups does not depend on income, then, 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. Selection into a reference group, however, is probably a function of income, with agents who have “rich” reference groups being more likely to observe higher-income individuals and vice-versa.

Fig. 1 illustrates the systematic biases that may arise with naïve agents. Fig. 1a depicts the income distribution for the whole population and for a rich reference group, which exhibits first-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 population). Since naïve agents use the information which they have 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. In Fig. 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. Fig. 1b illustrates this result by showing a middle-income reference group, where agents underestimate $F(y)$ for income levels $y < y^*$ and overestimate $F(y)$ for incomes $y > y^*$.

Agents with biased perceptions will obtain naïve estimates of many characteristics of the income distribution, such as the mean, median, dispersion and proportion of individuals under the poverty line, among others. For instance, if reference groups are more homogeneous in income than the total population (as is likely to be the case), perceptions about income inequality will be biased downward for all agents. This is consistent with Norton's and Ariely's (2011) finding that individuals systematically underestimate the level of inequality in the distribution of wealth in the United States.

A crucial parameter for this study is 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, it follows that they will underestimate their own rank in the distribution. Conversely, naïve agents with poor reference groups will overestimate their rank.

The main motivation for this discussion of biases is that misperceptions of the income distribution can have substantial implications for the determination of policy outcomes. This can be illustrated by incorporating biased perceptions into a basic framework like that of Meltzer and Richard (1981) with a simple redistributive scheme in which taxes on incomes above some cut-off point are used to finance benefits for agents with incomes below this level. If agents have biased perceptions of their own rank in the income distribution, their evaluations of how these costs and benefits will affect them are likely to be inaccurate. Naïve agents with poor reference groups will overestimate their own ranking in the overall distribution and may erroneously believe that they would not benefit from further income redistribution when they actually would. With access to the correct information about their actual place in the income distribution, self-interest would make these individuals change their attitude and favor, rather than oppose, the redistributive

² The base-rate neglect has been incorporated in economic models and empirical applications before (see Grether, 1990, and the reviews by Rabin, 1998, and DellaVigna, 2009).

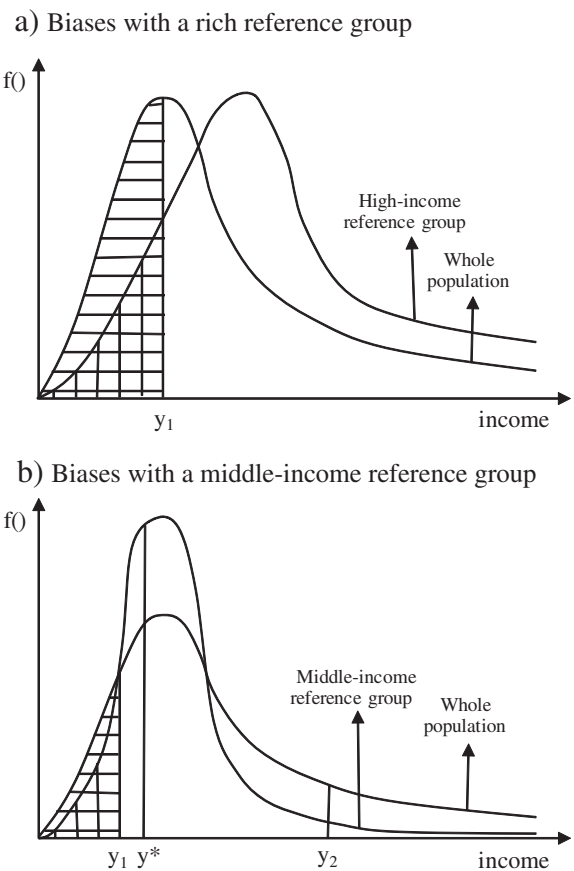


Fig. 1. Illustration of differences between population and reference-group income distributions.

policy. An analogous reasoning applies to individuals with rich reference groups. The experimental design of the survey used in this study allows for a direct test of this mechanism.

It should be noted that the recent political economy literature has discussed at length the fact that individuals may be motivated by more than self-interest, which implies that preferences for redistribution may reflect a wider set of factors.³ As a result, providing accurate information on the income distribution 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 identifies the net effect that providing accurate information on the income distribution will have on attitudes toward redistributive policies.

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 toward redistribution. The empirical research described in this paper is based on the *Survey on Distributional Perceptions and Redistribution*, a study of 1100 representative households representative in Greater Buenos Aires, Argentina. The survey was carried out in March 2009 and consisted of face-to-face interviews with a random sample of that population. It was specifically designed to test the mechanisms discussed in the previous section and, to that end,

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

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 for the cumulative subjective distribution: its evaluation at the point where each respondent thought his or her household stood.⁴ The question was worded as follows: "There are 10 million households in Argentina. Of those 10 million, how many do you think have an income lower than yours?"⁵ 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 monthly household levels. The intervals were chosen by the research team to correspond to the boundaries of deciles of the total national household income distribution at the time of the survey in order to facilitate the comparison of objective and perceived positions in the distribution using the experimental design.⁶

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 that was incorporated into the questionnaire. Randomized questionnaire-experiments have been developed in laboratory settings (Cowell and Cruces, 2004), while, in the context of household surveys and public opinion research, Horiuchi et al. (2007), Di Tella et al. (2012) and Hainmueller and Hiscox (2010) have conducted survey experiments with random allocations of questionnaire types among respondents.

As in these previous studies, the experimental setup for this survey involved randomly allocating two different types of questionnaires to interviewees, although the questions posed to the respondents were the same. The original feature of this setup has to do with the nature of the treatment, with the interviewer providing feedback to respondents in the treatment group in the form of accurate information concerning 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 households 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 households with a lower income than yours than you thought", (2) "You were right about how many households have a lower income than yours", or (3) "In fact, there are fewer households with a lower income than yours than you thought." 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 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 accurate information which differed from their stated perceptions.

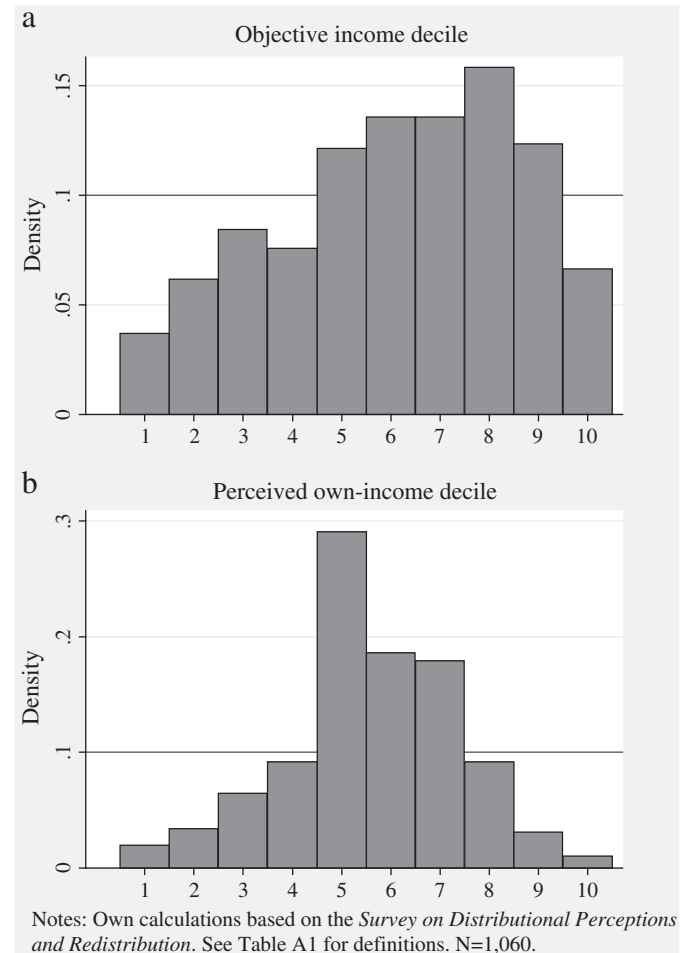


Fig. 2. Distribution of objective and perceived own-income decile.

⁴ 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 with answering in percentage terms.

⁵ 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).

⁶ The use of income intervals significantly reduces non-response rates. The notes provided in Table A1 (see the appendix) provide further details on the construction of the intervals and their implementation by interviewers in the survey.

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).

Source: Own calculations based on the Survey on Distributional Perceptions and Redistribution.

4. Evidence on perceptions of income distributions

4.1. Subjective income distributions

This section presents an analysis of the distribution of objective and perceived income rankings derived from the *Survey on Distributional Perceptions and Redistribution*. Fig. 2a gives 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 a 10% density. The higher concentration in deciles 5 to 9 is accounted for by the presence of higher average income levels in Buenos Aires relative to the countrywide average. Fig. 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 shown in Fig. 2a, the mode of the perceptions distribution is given by the fifth decile, with almost 30% 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 were thus substantially less dispersed than objective income levels are.

The difference between the two panels in Fig. 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. Table 1 presents summary statistics for this variable by objective income quintile. In all, 30% of respondents had positive biases, while 55% exhibited a negative bias: only about 15% of the respondents placed their household's income in the objectively correct decile. However, the deciles of a national income distribution are relatively narrow categories, and it is plausible that respondents could have difficulty in ascertaining their position with such precision. In any case, the proportion of biased individuals is more than 55% when considering only respondents who deviate from their true position by two deciles or more (see Cruces et al., 2011, for a lengthier discussion).

Table 1 also illustrates the relationship between the distributions of objective and perceived relative income levels depicted in Fig. 2. It is readily apparent that the average perceived own decile increases monotonically by quintile of objective income, although the range of the average perceived decile (from 4.60 for the bottom quintile to 6.48 for the top quintile) indicates that the distribution of perceptions is considerably more concentrated than that of objective income. This pattern has a direct correlate for the distribution of the bias in Table 1: respondents at the top and the bottom of the objective distribution display substantial negative and positive biases, respectively (of about −2.88 and 2.98 deciles for the extreme categories). Moreover, the bias diminishes up to the middle objective quintile, where it is close

to zero, and increases monotonically (in absolute values) from there onward. 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 for the whole population. If reference groups bundle together individuals of similar income levels, then one simple prediction 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 corresponds to the one depicted in Table 1. However, the observed pattern is also consistent with other potential explanations. For instance, poorer respondents may feel embarrassed to admit that their income is low and thus may over-report their true (accurate) perception, while richer individuals may not feel comfortable reporting their high relative position⁷ and thus may under-report their true (accurate) perception. The goal of this section is to test some predictions that are specific to the reference-group hypothesis. A further hypothesis is discussed in detail in the following subsection.

In the discussion presented in Section 2, the entire set of individual interactions (with friends, family, co-workers, etc.) was considered to constitute the relevant reference group for the formation of perceptions of income distribution. The analysis here uses a geographical proxy: the respondents' area of residence. Although it is not the best proxy for reference groups, the area of residence provides a simple illustration of 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).⁸ The survey covered 41 randomly selected sampling points (referred to as “neighborhoods” in the discussion) within 10 localities of the Greater Buenos Aires metropolitan area.⁹

The discussion in Section 2 pointed out that fully naïve agents will report their positions within their reference groups as their perceptions of their income ranks within the whole population; for “partially naïve” agents, relative income levels within the reference group will still be

⁷ 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.

⁸ The literature on interpersonal comparisons of well-being also proposes a geographic proxy for reference groups (Clark et al., 2008). The social networks literature also highlights the importance of area of residence for the exchange of information on employment and other income-generating activities (Bayer et al., 2008).

⁹ The sampling points correspond to a fairly small set of street blocks and contain 26 households, on average, in the sample. The average objective income level reported in the survey within each neighborhood ranges from an average objective decile of 3 to just below 8.

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 [0.0245]***	0.2099 [0.0280]***	−0.0168 [0.0944]	F-test [†]	0.0048 [0.1237]	0.0109 [0.1228]
Rank within locality	–	–	0.2151 [0.0868]**	0.2311 [0.1195]*	0.2002 [0.1114]*	0.2288 [0.1096]**
Has friends from all social classes	–	–	–	–	–	0.5046 [0.2897]*
Interaction: Locality rank & friends variable	–	–	–	–	–	−0.0937 [0.0458]**
Constant	4.0916 [0.1798]***	3.8997 [0.2266]***	4.2961 [0.2659]***	4.2846 [0.4379]***	4.1199 [0.6554]***	3.976 [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. The "rank within locality" variable was normalized to the same 1–10 scale used for the objective and subjective deciles. The "neighborhoods" correspond to 41 sampling points covering a small set of street blocks. These neighborhoods contain 26 households on average in the sample. The 10 "localities" represent larger geographical aggregates of the Greater Buenos Aires metropolitan area, with 106 observations on average.

Source: Own calculations based on the Survey on Distributional Perceptions and Redistribution.

partially correlated to those agents' perceptions of their income ranking. On the other hand, if individuals correctly apply Bayes' rule, relative income within a reference group should not have any explanatory power for perceived own-income rank within the population after controlling for the agent's objective overall position.

To illustrate the relevance of the reference-group hypothesis, Table 2 presents a series of regressions where the perceived own-income decile is the dependent variable.¹⁰ Column 1 presents a simple regression with the respondents' objective income deciles as the sole independent variable. The results in this column confirm the existence of a highly significant relationship between the two variables discussed above, with the coefficient being positive. This coefficient falls short of 1 partly because of systematic biases in perceptions, but also because of the attenuation bias stemming from the presence of measurement error in the independent variable. The second column repeats this simple regression, but includes 41 neighborhood fixed effects: the coefficient is still significant at the 1% level, although it is slightly lower.¹¹

The first 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), converted to the same 1–10 scale used for the objective and subjective deciles. 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 coefficient of 0.2151. The respondents' relative incomes within their localities thus seem to have a strong correlation with their perceptions of the distribution, even after controlling for their objective income levels. 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 indicator variables. 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 cannot reject the hypothesis that the objective income indicators are all equal to zero (*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 notes included in Table 2 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.

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 an indicator variable that takes the value 1 if an individual answered that his or her friends are from all social classes (38.5% of the respondents) and zero otherwise. Intuitively, people who interact with several distinct groups must have more information about the role of income in forming reference groups (or, alternatively, the selection process is more salient for them, so they are less likely to fail to consider Bayes' rule). As a consequence, they should be less inclined 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 this indicator variable, as well as its interaction with the respondents' income rank within the locality. The coefficient of the interaction is negative and significant, as expected: conditional on their own objective income level, individuals with broader reference groups should be less influenced by their relative income within their neighborhood. However, the coefficient of the variable without the interaction is also significant, indicating some correlation between income rank and this variable.

¹⁰ 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.

¹¹ Standard errors are clustered at the neighborhood level. All the results in the table are similar if 10 locality fixed effects (with robust standard errors) are included instead.



Note: Based on the auxiliary February 2012 household survey, Greater La Plata, Argentina (see Section 4.3 for details). Answer to the question: “How much confidence do you have about the reply you just gave?” referring to the previous question on the individual’s perceived own-income decile. The options were “Not sure at all”, “Somewhat sure”, “Sure” and “Very sure” (1 to 4). The sample mean is represented by the solid horizontal line. 95 percent confidence intervals computed for each decile. N=302.

Fig. 3. Confidence in response, 1 (low) to 4 (high) scale, by perceived own-income decile.

4.3. Alternative explanations

Section 4.2 presented suggestive evidence that the reference-group conjecture presented in Section 2 can explain a substantial variation in the observed distribution of biases. However, other competing or complementary explanations may also account for the observed patterns. A powerful alternative hypothesis is the prevalence of focal-point answers. 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. 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 provides several examples of survey responses with focal points at 50 for distributions of between 0 and 100.

In this study, anchoring to the middle of the scale would be a concern if it is induced by epistemic uncertainty, which implies that individuals use the “50%” response when assessing probabilities in cases where, in fact, they do not know the answer to the question (Bruine de Bruin et al., 2002). Bruine de Bruin and Carman (2012) find strong evidence in support of this hypothesis in the context of questions about the probability of own-survival for adults. This would translate into a clustering of answers in the fifth (or sixth) decile and would generate patterns similar to those documented in Section 4.1 (e.g., rich people apparently tend to underestimate their own income ranking, while the poor tend to do the opposite). However, these biases could be an artifact created by the respondents’ lack of knowledge rather than by the influence of reference groups.

A second auxiliary survey was conducted to provide a formal means of testing whether epistemic uncertainty plays a role in the biases documented in Section 4.2. Power calculations indicated that a smaller sample was needed, and a final sample of 302 completed questionnaires was therefore used.¹² The questionnaire included a small set of background questions on such matters as age, gender, education and income level (proxied by decile), as well as the own-income-rank evaluation question. Building on Bruine de Bruin and Carman (2012), immediately after the rank perception question, the respondent was asked the following: “You just told me that X million households have a lower

income than yours. How sure are you of the answer that you just gave?” There were four possible options for a response to this question: “Not sure at all”, selected by 8% of the respondents; “Somewhat sure” (48%); “Sure” (37%); and “Very sure” (7%). The average for the variable (coded as 1 to 4) is 2.43, with a standard deviation of 0.74.

As depicted in Fig. 3, the average of the confidence variable (coded as 1 to 4) is not significantly different for those placing themselves at deciles 5 or 6 and, in fact, the confidence intervals are consistent with a roughly constant level of confidence in responses across the entire distribution of perceived deciles. The hypothesis of no differences in certainty between respondents with a perceived own-decile equal to 5 and the rest cannot be rejected for conventional significance levels. The same pattern of results remains with alternative definitions of certainty (e.g., indicators for “Not sure at all”, or for “Not sure at all” and “Somewhat sure”), as well as when comparing those who selected deciles 5 or 6 with the rest of the respondents (results not reported). In summary, epistemic uncertainty does not seem to be a predominant force underlying the systematic biases in perceptions of own-income rank documented above. This might be due to the fact that the own-income-rank evaluation question was phrased in terms of “millions of households” having lower incomes than the respondent’s – that is, as a question about frequencies rather than one designed to elicit a reply given in terms of percentages or probabilities (see the discussion of frequency representations and biases in Mellers et al., 2001).

5. Biased perceptions and preferences for redistribution: experimental results

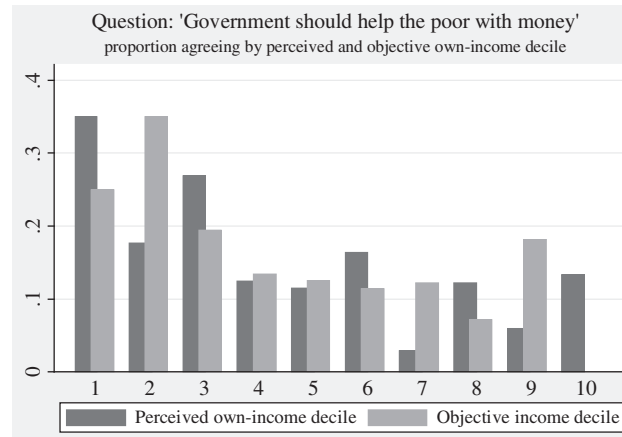
5.1. Identification strategy

As described in Section 3, the survey included a field questionnaire-experiment: the interviewer provided a randomly assigned group of respondents with unbiased estimates of their positions in the income distribution, pointing out the degree and direction of the bias in each respondent’s self-assessment (if any). This section discusses the causal effect of this information treatment on preferences for redistribution.¹³

Let the expression $r_i = F_i(x_i)$ represent the (potentially biased) perception of an agent’s own income rank in the income distribution. Let $d(r_i)$ be a function that maps out the relationship between an individual’s perceived rank and his or her demand for redistribution. For example, $d(\cdot)$ can represent a binary variable denoting support

¹² The survey design is identical to that of the *Survey on Distributional Perceptions and Redistribution*. It was implemented on February 11 and 12, 2012. It consisted of 302 face-to-face interviews of people in a random sampling based on 24 sampling points in the Greater La Plata metropolitan area (the capital of the Province of Buenos Aires and its surroundings). This area borders the Greater Buenos Aires metropolitan area, where the original survey was carried out, and has a similar socioeconomic profile.

¹³ This discussion has its origins in a suggestion by an anonymous referee.



Notes: Own calculations based on the *Survey on Distributional Perceptions and Redistribution*. See Table A1 for definitions. Observations for individuals in the control group (i.e., individuals that were not provided feedback about the accuracy of their income ranking assessments). N=512.

Fig. 4. Preferences for redistribution – agreement with “help the poor with money” question, by objective and perceived own-income decile (control group only).

for a program that taxes households above a given quantile and redistributes it to individuals below that quantile. In the simple Meltzer and Richard (1981) framework discussed in Section 2, where agents are purely selfish, $d(\cdot)$ would be 1 for individuals below the cutoff quantile and zero for those above the cutoff.

However, preferences for redistribution are not necessarily that simple. The shape of the function $d(\cdot)$ depends on the underlying individual preferences for redistribution in the population under study. While $d(\cdot)$ may be monotonically decreasing in income if only self-interest is at work, several of the factors discussed in Section 2 indicate that this relationship might not be monotonic in income. For instance, a U-shaped function $d(\cdot)$ would denote a situation where a redistributive program is supported by those with very high and very low income levels.

The informational treatment in the experimental setup can be assumed to affect individual preferences for redistribution through its effect on r_i , the perceived own-income rank. Expressing rankings in terms of deciles, when an individual who thinks that he or she is in decile r_i is informed that his/her position is actually in decile q_i , the resulting change in that individual's perception could be modeled by an update function $u(\cdot)$: the new perception is $r_i + u(r_i, q_i)$. If $r_i > q_i$ (if the individual initially perceived his/her income rank as higher than it actually is), then $u(r_i, q_i)$ is expected to be negative – i.e., the individual updates his/her beliefs downward. Since the informational treatment is relatively weak, consisting of a verbal remark made by the interviewer,¹⁴ it is likely that there will be no more than a partial updating of perceptions: $|u(r_i, q_i)| < |r_i - q_i|$. This condition implies also that $u(r_i, q_i) = 0$ if $r_i = q_i$.

Providing individuals with objective information about their income rank (q_i) has the following effect on support for redistribution: $\Delta\text{Support} = d(r_i + u(r_i, q_i)) - d(r_i)$. The identification of this effect depends on several factors. A first issue is that, given a pair (r_i, q_i) , under the null hypothesis of a zero-treatment effect, it is not possible to determine whether the absence of a significant effect is due to the fact that the provision of information did not affect the individual's perception, $u(r_i, q_i) = 0$; whether it changed the perception but the support was not affected by this change, i.e. $d(r_i + u(r_i, q_i)) = d(r_i)$ even though $u(r_i, q_i) > 0$; or both.¹⁵

¹⁴ 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 weak treatment.

¹⁵ For example, if an individual's level of support is the same with a perceived own-income decile equal to 7 and 8, and that individual perceives his or her position to be at 8 when the income level actually corresponds to decile 7, providing that individual with unbiased information will not affect his/her support for redistribution even if it successfully changes the individual's perception of his/her relative income.

A second issue is that the treatment effect can be very heterogeneous, depending on the shape of both functions $d(\cdot)$ and $u(\cdot)$. An ideal empirical setup would estimate the effect of the informational treatment for each separate pair (r_i, q_i) . For instance, the experimental design would randomly provide information within the group of individuals with a perceived own-income decile equal to 5 and an objective income decile equal to 2, and so forth for every possible pair (r_i, q_i) . However, this estimation would demand a very large sample size to ensure a critical number of individuals in each cell (r_i, q_i) .

In practice, since the sample size is limited, the empirical strategy must rely on the estimation of treatment effects for coarser subgroups of the population. The extreme case involves assuming that the effect of telling an individual that his/her perception of own-income rank is biased by X deciles is symmetric: e.g., the effect of telling an individual with a perceived own-income decile equal to 5 that he/she is actually in decile 4 has the same magnitude (but opposite sign) as the effect of telling an individual with a perceived own-income decile equal to 5 that he/she is in fact in decile 6.¹⁶ However, there are important reasons to believe that the treatment effect is not symmetric for those with positive and negative biases.

First, the indicators of support for redistribution discussed below do not appear to be monotonic in perceived rank, implying that the shape of the function $d(\cdot)$ follows a more complex pattern. Fig. 4 depicts the average agreement with a question as to whether the government should help the poor with monetary subsidies by level of objective and perceived own-income decile for the control group in the experiment (those who did not receive feedback from the interviewer on their actual position in the distribution). While the average of this variable is higher for the three lowest perceived own-income deciles, this relationship does not appear to be monotone. More precisely, the findings suggest that informing individuals who perceive themselves as being in the middle of the distribution that they actually are among the poorest would have a substantially greater effect on preferences for redistribution (in absolute values) than the effect on those preferences that would be generated by informing individuals who place themselves in the middle of the distribution that they are richer than they thought.

¹⁶ Another simplifying assumption might be that the effect is linear: for instance, the effect of informing an individual with a perceived own-income decile equal to 5 that he/she is actually in decile 4 should be half as great as the effect of correcting the perception by two deciles (i.e., from a perception of an own-income decile of 5 to an objective decile of 3).

A second source of asymmetry in the treatment effect for those with positive and negative biases is that the function $u(\cdot)$ may differ substantially between those two groups of individuals. Those with negative biases tend to be in the upper half of the income distribution and therefore tend to have not only higher incomes but also higher levels of education. Richer and more educated individuals may be better informed and thus more confident about their beliefs (i.e., a Bayesian prior with lower variance), which would make them less reactive to the informational treatment. This intuition is supported by data from the auxiliary survey described in Section 4.3. Fig. 3 indicates that those in the top perceived own-income decile appear to be more certain about their assessment of their position in the income distribution than those in the bottom decile. Moreover, more educated individuals report significantly higher levels of confidence in their assessment of their own relative rank (results not reported).

Given these potential asymmetries, the most natural choice is to estimate the effect of the informational treatment for three groups: those with a positive bias, those with a negative bias, and those without any bias. For respondents with no bias, the treatment simply confirms their perception of their own rank in the distribution. Since they do not receive any new information, the treatment should be immaterial to their stated preferences for redistribution. The existence of an impact of the treatment for this group could signal that the interviewer's statement had an effect that was independent of its content, which could be difficult to separate from that of the actual information provided.¹⁷ The treatment ought to prompt those who overestimated their own rank to demand more help for the poor, insofar as it makes them more likely to consider themselves as potential beneficiaries of this type of policy. Finally, for those who underestimated their rank, the treatment can be expected to make them less likely to believe that they will benefit from the program, so self-interest should induce them to demand less redistribution (although the non-monotonicity of $d(\cdot)$ may prevent this from happening).

5.2. Experimental results

The focus of this study is on attitudes toward specific policies designed to help the poor rather than on general beliefs about justice and income redistribution. The questions used in this section are fairly general, but they refer to specific ways in which government programs to help the poor should be implemented. Direct government transfers to the poor were especially relevant in the context of the population under study. 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 that 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 direct cash transfers to poor families, which were ultimately introduced 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. In addition to their pivotal role in the political debate of interest to the population under study, another advantage of studying preferences regarding broadly defined policies rather than broadly defined distributive principles is that government cash transfer programs can reach a substantial proportion of the population and can be fairly accurately targeted at the

¹⁷ 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 is suggestive that under-reporting is not the driving force behind the revealed biases.

poor, unlike most of Argentina's other items of public expenditure (Gasparini and Cruces, 2010). This implies that low-income respondents faced a significant probability of being directly affected by policy changes in this area.¹⁸

Individuals in the treatment group were given unbiased information about their income ranking after they had reported their own perception. The post-treatment questions¹⁹ 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 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% of the interviewees (ranging from 23.2 in the bottom quintile to 9.5% 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% of respondents, with the percentage also decreasing monotonically by income level (42.1% to 22.1% 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% of the respondents, with a virtually constant proportion across income groups. While the low level of variation in this variable implies that it cannot be studied independently, it is included in the composite measures of support for redistribution.

The correlation between the answers to the questions about money and food is 0.44 and significant at the 1% level, which indicates that the joint analysis of these responses in a composite variable is warranted.²⁰ The main dependent variable for the analysis consists of a composite indicator equal to the average response to the three forms of government assistance listed above. This has an average of 0.49 (ranging from 0.54 in the poorest quintile to 0.43 in the richest quintile). The results are also reported in terms of an alternative composite measure, a standardized index which removes the means and weights of each of the three measures by their standard deviations (Kling et al., 2007). This measure, reported in effect sizes, has a mean of zero and ranges from 0.134 in the poorest quintile to -0.117 in the richest quintile.

The dependent variables in Table 3 are the different measures of preferences for redistribution (*mean support*, *standardized mean*, *help with money* and *help with food*) and a fifth dependent variable, which equals 1 if the respondent reports having made a donation in the past 12 months (this is used as a falsification test, as discussed below). The results are presented for each of the treatment subgroups: those with a negative bias, those with a zero bias, and those with a positive bias. The first two rows in each panel display the means for each dependent variable by treatment status. The following row shows the difference between the two and the standard error of this difference, with stars denoting the significance of the mean difference test. Finally, the last row in each panel presents the results from the regression version of the test, i.e., the coefficient of the treatment indicator in an OLS regression that includes the same control

¹⁸ A well-targeted emergency program was implemented in 2002. It covered the poorest 20% of households in the country. The new cash transfer program (launched in 2009) reached a similar level of coverage and effectiveness in targeting in 2010, a year after its implementation (Cruces and Gasparini, 2010).

¹⁹ Table A2 (see 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.

²⁰ There are several plausible explanations 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).

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

	Negative bias: Treatment = telling respondents that position is higher	No bias: Treatment = confirming respondent's positional perception	Positive bias: Treatment = telling respondent that position is lower	Positive bias: More than one decile
	(1)	(2)	(3)	(4)
<i>Mean of three government-support-to-the-poor questions (money, food, jobs)</i>				
Treatment group [obs.]	0.459 [296]	0.532 [84]	0.538 [150]	0.559 [99]
Control group [obs.]	0.463 [286]	0.495 [72]	0.509 [152]	0.500 [112]
Difference [s.e.]	−0.003 [0.018]	0.036 [0.041]	0.029 [0.029]	0.059 [0.034]**
Conditional diff. [s.e.]	−0.003 [0.015]	0.015 [0.066]	0.071 [0.026]***	0.096 [0.042]**
<i>Standardized mean (Kling et al., 2007), support questions (money, food, jobs)</i>				
Treatment group [obs.]	−0.067 [296]	0.126 [84]	0.109 [150]	0.192 [99]
Control group [obs.]	−0.063 [286]	0.026 [72]	0.048 [152]	0.009 [112]
Difference [s.e.]	−0.004 [0.049]	0.101 [0.108]	0.060 [0.084]	0.182 [0.094]**
Conditional diff. [s.e.]	−0.004 [0.045]	0.035 [0.183]	0.179 [0.082]**	0.259 [0.116]**
<i>Government should help the poor with money</i>				
Treatment group [obs.]	0.111 [296]	0.226 [84]	0.212 [151]	0.242 [99]
Control group [obs.]	0.108 [287]	0.153 [72]	0.176 [153]	0.150 [113]
Difference [s.e.]	0.003 [0.026]	0.073 [0.063]	0.035 [0.045]	0.092 [0.054]**
Conditional diff. [s.e.]	0.010 [0.021]	0.063 [0.119]	0.084 [0.046]*	0.132 [0.067]*
<i>Government should help the poor with food</i>				
Treatment group [obs.]	0.284 [296]	0.381 [84]	0.424 [151]	0.444 [99]
Control group [obs.]	0.303 [287]	0.347 [72]	0.373 [153]	0.381 [113]
Difference [s.e.]	−0.019 [0.038]	0.034 [0.078]	0.051 [0.056]	0.064 [0.068]
Conditional diff. [s.e.]	−0.024 [0.034]	−0.010 [0.105]	0.118 [0.046]**	0.124 [0.071]*
<i>Falsification test: Made donations in the last twelve months</i>				
Treatment group [obs.]	0.866 [292]	0.788 [85]	0.719 [153]	0.687 [99]
Control group [obs.]	0.831 [284]	0.845 [71]	0.742 [155]	0.722 [115]
Difference [s.e.]	0.035 [0.030]	−0.057 [0.063]	−0.023 [0.051]	−0.035 [0.063]
Conditional diff. [s.e.]	0.044 [0.030]	−0.024 [0.082]	0.029 [0.054]	0.031 [0.064]

Notes: * represents statistical significance at the 10% level; ** at the 5% level; and * at the 1% level. These levels correspond to the test for the unconditional difference – 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 3–4). 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 significance levels underlying the stars are derived from standard errors clustered at the neighborhood level. 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 own income decile minus objective income decile. See Table A1 for further variable definitions.

Source: Own calculations based on the Survey on Distributional Perceptions and Redistribution.

variables from columns 5 and 6 of Table 2 (standard errors are also clustered at the neighborhood level).²¹

The results shown in the second column of Table 3 indicate that some of the point estimates of the differences between treatment and controls for those with a zero bias are non-negligible. For instance, the difference for the mean of the three variables (first panel) is 0.036. However, none of the differences (conditional or unconditional) for any of the dependent variables that were considered are statistically different from zero. These results are compatible with the discussion in the previous section, which indicated that the treatment could be expected to have no effect on preferences for redistribution for this group.

The first column of Table 3 presents the treatment effect for individuals with negative biases. The differences between the treatment and control groups are relatively small for all the dependent variables considered (positive for *help with money* and negative for *help with food*, and very close to zero for the two aggregate variables), with the exception of the donations variable, which exhibits slightly higher differences. None of the point estimates, however, are statistically significant at conventional levels.

Finally, the results for individuals with positive biases (those who overestimate their rank), shown in the third column, point to a series of relatively large and statistically significant differences between treatments and control groups. As predicted, treated individuals in this

group exhibited higher levels for the indices of support for redistribution. For instance, the unconditional (conditional) difference is 0.029 (0.071) for the *mean support* variable.²² The p -values of the unconditional differences in the four outcomes of interest vary from 0.162 (*mean support*) to 0.236 (*standardized mean*), but the conditional differences for the four variables are in all cases statistically significant at conventional levels for this group (at the 10% level for *help with money*, at the 5% level for *help with food* and the *standardized mean*, and at the 1% level for *mean support*).

In line with the discussion presented in the previous subsection, the effect of the information treatment should be a function of the degree of bias in perceptions of the income distribution: 30.6% of those with a positive bias had a misperception of only one decile, and thus the informational treatment does not necessarily convey a strong message. The fourth column reports the treatment effects for the subsample of individuals who overestimated their own ranking by more than one decile.²³ Since this group exhibits a substantially larger gap

²² The main goal of using conditioning variables in an experimental setting is to improve efficiency. If the treatment assignment is randomized, *in expectation*, the inclusion of control variables should not change much the point estimates. Following Altonji et al. (2005), if the correlation between the treatment and the observables is informative about the correlation between the treatment and the unobservables, then the larger coefficients obtained when including controls indicate that the conditional difference under-estimates the true effect of the treatment. If the coefficients had decreased instead, the potential problem would be the over-estimation of the causal effect of the treatment.

²³ The analogous situation for the group considered in column one of the table would refer to those with a negative bias of two deciles or more. None of the differences (conditional or unconditional) in the preferences for redistribution variables are statistically significant at the standard levels for this group. These results are omitted due to space constraints.

²¹ The inclusion of control variables can reduce the variability of the error term, which increases the statistical power. This is an important adjustment in this context due to the small sample size within each group.

between perceptions and the actual situation, the effect on support for redistribution can be expected to be stronger.²⁴ The evidence set forth in Table 3 indicates that this is indeed the case: the treatment effects shown in the fourth column are substantially higher than those shown in the third column and are statistically significant even for the unconditional estimates. The differences in support for redistribution between the treatment and control groups are large and statistically significant for the two indices (both conditional and unconditional) and for the two components (the conditional and unconditional difference for *help with money* and the conditional difference for *help with food*).²⁵

Finally, the bottom panel in Table 3 presents a falsification test designed to capture the presence of any spurious effects of the treatment on respondents. These are estimates of the effect of the informational treatment on the variable defined by the post-treatment question: “Have you made any donations to an individual or charity during the past 12 months?” This donation variable was included in the survey because of its close relationship to a willingness to provide assistance to the poor. If the treatment has an effect through a change in the perceptions of own-income rank, it should have an impact on stated preferences, but should not affect the reporting of past actions. If, on the contrary, the provision of information induces more caring or generous statements from the respondent through a shaming effect, then the treatment should 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, despite some sizeable differences between treatments and controls for some of the groups, none of these differences (conditional or unconditional) are statistically significant at the standard levels.

To sum up, the evidence for those with no bias in their perceptions supports the discussion in the previous section, which indicated that the treatment should not have an impact on preferences for redistribution for individuals in this group. Moreover, the evidence also suggests that the treatment did not have a systematic effect for individuals with negative biases. One plausible explanation for that finding is that the treatment affected perceptions but failed to impact attitudes toward redistribution because $d(\cdot)$ is roughly flat for the relevant range of income ranks. Alternatively, the informational treatment may not change the perceptions of own-income rank for these individuals (i.e., $u(r_i, q_i) = 0$). This may have happened because respondents did not find the information credible, or simply because they had a strong prior belief. Consistent with this explanation, richer and more educated individuals have greater confidence in their assessment of their own relative rank, as discussed in Section 4.3.

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 income rank. Moreover, these effects are larger and systematically significant for those with biases greater than one, that is, for those with higher degrees of misperception. These effects 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 individuals in the control group.

6. Conclusion

The motivation for undertaking this study was the apparent lack of evidence regarding 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 focuses on a simple mechanism, whereby agents extrapolate from their reference group without accounting for the selection process underlying the formation of the group owing to either informational or cognitive limitations. A tailored household survey provided evidence of the presence of sizeable systematic biases in perceptions about relative income in Argentina. Furthermore, the analysis uncovered suggestive evidence that the incorrect extrapolation of information from reference groups can be a powerful explanation for the observed biases. The systematic biases documented in this paper and their consequences can arise in any society.

This study involved the implementation of an original survey experiment in the field in which a randomly assigned group of respondents was provided with accurate information about their ranking in the income distribution as a form of feedback concerning their responses. Confronting agents who had biased perceptions with this information had a significant effect on their stated preferences for redistribution: those who overestimated their relative position (who thought that they were 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.²⁶ Having accurate information about the income distribution might induce agents to better calibrate their demands for redistribution. The 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 for improving policies, although the impact of the biases in the efficiency of redistribution should also be considered (Acemoglu and Robinson, 2001).

The role of misconceptions in political economy has been studied before (Romer, 2003; Slemrod, 2006). While Besley (2007) highlights the benefits to be derived from incorporating notions of dispersed and limited information for modern political economy, building-in more specific factors, such as 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 those pointed out by 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 how perceptions may affect stated preferences for redistribution. Further research could focus on the impact of biases and information on actual behavior, such as voting patterns. Moreover, it would be interesting to explore how misperceptions affect individuals' reactions to redistributive policies (e.g., as expressed through charitable contributions and tax evasion) and to learn whether the provision of information on broader issues that go beyond such matters as rules and regulations (e.g., Chetty and Saez, 2009) may have implications for public finance models. Finally, the results of the analyses conducted in connection with this study could be attributable either to limited information or to limited cognitive ability — further research will be needed in order to pinpoint the source of the observed biases in distributional perceptions.

²⁴ The same intuition applies for other groups for which the information treatment indicated a large discrepancy between their perception and the actual situation. Cruces et al. (2011) report further robustness tests of this kind.

²⁵ However, as expected from the discussion of asymmetric reactions to the informational treatment, there are no significant differences in support for redistribution between the treatment and control groups among those with negative biases greater than 1 (in absolute values).

²⁶ Since the bias affects preferences for redistribution for relatively poorer individuals, it is reminiscent of the Marxian notion of false consciousness. 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.”

Appendix A

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>I will show you a table with levels of income. Please indicate where, approximately, you would say that your household is located, considering all income in your household from every source (work, government transfer programs, pensions, rent, etc.) (1) Less than X; (2) X to Y; ...; (10) More than Z</i>	6.12	2.46	1	10	1060
Perception of own income decile	The interviewer made a statement and asked the following question: <i>There are 10 million households 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 million; ...; (10) 9 to 10 million.</i>	5.60	1.77	1	10	1060
Bias	The bias is constructed as the level of the <i>objective income decile</i> minus that of the <i>perceived own-income decile</i> . 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.	-0.53	2.49	-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 <i>objective income decile (X)</i> to that of the <i>perception of own income decile (Y)</i> , and stating accordingly: The interviewer read the following statement (with X and Y being determined by previous answers): <i>"Based on your income level, the latest studies conducted by the University indicate that there are X million households with a lower income than yours, while you stated that there were Y."</i> The interviewer then read out one of the three following statements, depending on the accuracy of the X/Y comparison: (1) <i>"In fact, there are more households with a lower income than yours than you believed"</i> , (2) <i>"You were right about how many households have a lower income than yours"</i> , or (3) <i>"In fact, there are fewer households with a lower income than yours than you thought"</i> .	0.514	0.500	0	1	1060
Rank within locality	This variable is constructed using the <i>objective income decile</i> variable for each respondent and computing his/her ranking within his/her area of residence, where the 41 sampling points were 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 a lower income than that of the respondent divided by the total number of households.	5.42	2.48	1.08	9.78	1060
Respondent has friends from all social classes	This is an indicator variable equal to 1 if the respondent states that he or she has friends from all social classes when asked the question: <i>"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)?"</i>	0.376	0.485	0	1	1060
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
Support for redistribution: Mean	Variable representing the mean of the response to the three previous questions.	0.487	0.237	0	1	1040
Support for redistribution: Standardized mean	Aggregate variable based on the <i>help with money</i> , <i>help with food</i> and <i>help find jobs</i> questions. The process consists of demeaning each variable, converting each to effect sizes by dividing the result by the control group's standard deviation, and taking the average of the effect sizes (Kling et al., 2007)	-0.002	0.6509	-2.87	1.35	1040
Respondent made donations in the past 12 months	"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) <i>Primary incomplete</i> ; (2) <i>Primary complete</i> ; (3) <i>Secondary incomplete</i> ; (4) <i>Secondary complete</i> ; (5) <i>Undergraduate incomplete</i> ; (6) <i>Undergraduate complete</i> ; (7) <i>Postgraduate</i> . 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

Notes: 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.

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. The boundaries of the intervals correspond to actual deciles of the distribution, which facilitated the comparison of objective and perceived rank as 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% interval of the actual points in the microdata.

The two versions of the questionnaire (and their English translations) will be included in an online appendix.

Source: Own calculations based on the Survey on Distributional Perceptions and Redistribution.

Table A2

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

Variable	Treatment group	Control group	Difference ($\mu_t - \mu_c$)	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)	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
Informal 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 1115 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% 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.

Source: Own calculations based on the Survey on Distributional Perceptions and Redistribution.

References

- Acemoglu, D., Robinson, J., 2001. Inefficient redistribution. *American Political Science Review* 95 (3), 649–661.
- Alesina, A., Angeletos, G.M., 2005. Fairness and redistribution: US vs. Europe. *The American Economic Review* 95, 913–935.
- Alesina, A., Giuliano, P., 2009. Preferences for redistribution. NBER Working Paper No. 14825.
- Alesina, A., Glaeser, E., 2004. *Fighting Poverty in the US and Europe. A World of Differences*. Oxford University Press.
- Alesina, A., La Ferrara, E., 2005. Preferences for redistribution in the land of opportunities. *Journal of Public Economics* 89 (1), 897–931.
- Altonji, J., Elder, T., Taber, C., 2005. Selection on observed and unobserved variables: assessing the effectiveness of catholic schools. *Journal of Political Economy* 113 (1), 151–184.
- Bartels, L., 2008. *Unequal Democracy: the Political Economy of the New Gilded Age*. Russell Sage Foundation and Princeton University Press.
- Bayer, P., Ross, S., Topa, G., 2008. Place of work and place of residence: informal hiring networks and labor market outcomes. *Journal of Political Economy* 116 (6), 1150–1196.
- Bénabou, R., Ok, E., 2001. Social mobility and the demand for redistribution: the POUM hypothesis. *Quarterly Journal of Economics* 116 (2), 447–487.
- Bénabou, R., Tirole, J., 2006. Belief in a just world and redistributive politics. *Quarterly Journal of Economics* 121 (2), 699–746.
- Benoît, J.P., Dubra, J., 2011. Apparent overconfidence. *Econometrica* 79 (5), 1591–1625.
- Besley, T., 2007. The new political economy. *The Economic Journal* 117 (524).

- Bruine de Bruin, W., Carman, C.G., 2012. Measuring risk perceptions: what does the excessive use of 50% mean? *Medical Decision Making* 32 (2), 232–236.
- Bruine de Bruin, W., Fischbeck, P.S., Stiber, N.A., Fischhoff, B., 2002. What BNumber is 'fifty-fifty'? Redistributing excess 50% responses in risk perception studies. *Risk Analysis* 22, 725–735.
- Card, D., Mas, A., Moretti, E., Saez, E., 2010. Inequality at work: the effect of peer salaries on job satisfaction. NBER Working Paper No. 16396.
- Chetty, R., Saez, E., 2009. Teaching the tax code: earnings responses to an experiment with EITC claimants. NBER Working Paper No. 14836.
- Clark, A., Frijters, P., Shields, M., 2008. Relative income, happiness, and utility: an explanation for the Easterlin paradox and other puzzles. *Journal of Economic Literature* 46 (1), 95–144.
- Cowell, F., Cruces, G., 2004. Perceptions of inequality and risk. *Research on Economic Inequality* 12, 99–132.
- Cruces, G., Gasparini, L., 2010. Las asignaciones universales por hijo en Argentina: impacto, discusión y alternativas. *Economica* 56 (1), 105–146.
- Cruces, G., Pérez-Truglia, R., Tetaz, M., 2011. Biased perceptions of income distribution and preferences for redistribution: evidence from a survey experiment. IZA Working Paper 5699.
- DellaVigna, S., 2009. Psychology and economics: evidence from the field. *Journal of Economic Literature* 47, 315–372.
- Di Tella, R., Galiani, S., Schargrodsky, E., 2012. Reality versus propaganda in the formation of beliefs about privatization. *Journal of Public Economics, Elsevier* 96 (5), 553–567.
- Duflo, E., Saez, E., 2003. The role of information and social interactions in retirement plan decisions: evidence from a randomized experiment. *Quarterly Journal of Economics* 118 (3), 815–842.
- Fong, C., 2001. Social preferences, self interest and the demand for redistribution. *Journal of Public Economics* 82 (2), 225–246.
- Gasparini, L., Cruces, G., 2010. A distribution in motion: the case of Argentina. In: López Calva, Luis Felipe, Lustig, Nora (Eds.), *Declining Inequality in Latin America: A Decade of Progress?* Brookings Institution Press.
- Glaeser, E., Resseger, M., Tobio, K., 2008. Urban inequality. NBER Working Paper No. 14419.
- Grether, D.M., 1990. Testing Bayes rule and the representativeness heuristic: some experimental evidence. *Journal of Economic Behavior and Organization* 17, 31–57.
- Hainmueller, J., Hiscox, M., 2010. Attitudes toward highly skilled and low-skilled immigration: evidence from a survey experiment. *American Political Science Review* 104, 61–84.
- Horiuchi, Y., Imai, K., Taniguchi, N., 2007. Designing and analyzing randomized experiments: application to a Japanese election survey experiment. *American Journal of Political Science* 51 (3), 669–687.
- Hurd, M., 2009. Subjective probabilities in household surveys. *Annual Review of Economics* 1, 543–564.
- Jensen, R., 2010. The perceived returns to education and the demand for schooling. *Quarterly Journal of Economics* 125 (2), 515–548.
- Kahneman, D., Tversky, A., 1972. Subjective probability: a judgment of representativeness. *Cognitive Psychology* 3, 430–451.
- Kahneman, D., Slovic, P., Tversky, A. (Eds.), 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press.
- Kling, J.R., Liebman, J.B., Katz, L., 2007. Experimental analysis of neighborhood effects. *Econometrica* 75 (1), 83–119.
- Luttmer, E., 2001. Group loyalty and the taste for redistribution. *Journal of Political Economy* 109 (3), 500–528.
- Manski, C., 2004. Measuring expectations. *Econometrica* 72 (5), 1329–1376.
- Mellers, B.A., Hertwig, R., Kahneman, D., 2001. Do frequency representations eliminate conjunction effects? An exercise in adversarial collaboration. *Psychological Science* 12, 269–275.
- Meltzer, A., Richard, S., 1981. A rational theory of the size of government. *Journal of Political Economy* 89 (5), 914–927.
- Norton, M., Ariely, D., 2011. Building a better America—one wealth quintile at a time. *Perspectives on Psychological Science* 6, 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* 93 (7–8), 950–964.
- Piketty, T., 1995. Social mobility and redistributive politics. *Quarterly Journal of Economics* 110 (3), 551–584.
- Rabin, M., 1998. Psychology and economics. *Journal of Economic Literature* 36 (1), 11–46.
- Ravallion, M., Lokshin, M., 2002. Self-rated economic welfare in Russia. *European Economic Review* 46 (8), 1453–1473.
- Romer, P., 2003. Misconceptions and political outcomes. *The Economic Journal* 113 (484), 1–20.
- Rotemberg, J., 2002. Perceptions of equity and the distribution of income. *Journal of Labor Economics* 20 (2) (part 1).
- Simon, H., 1972. Theories of bounded rationality chapter 8. In: McGuire, C.B., Radner, R. (Eds.), *Decision and Organization*. North-Holland Publishing Company, Amsterdam.
- Slemrod, J., 2006. The role of misconceptions in support for regressive tax reform. *National Tax Journal* 59, 57–75.
- Zafar, B., 2011. How do college students form expectations? *Journal of Labor Economics* 29 (2), 301–348.