

Why Are Relatively Poor People Not More Supportive of Redistribution? Evidence from a Randomized Survey Experiment across Ten Countries[†]

By CHRISTOPHER HOY AND FRANZISKA MAGER*

We test a key assumption underlying seminal theories about preferences for redistribution, which is that relatively poor people should be the most in favor of redistribution. We conduct a randomized survey experiment with over 30,000 participants across 10 countries, half of whom are informed of their position in the national income distribution. Contrary to prevailing wisdom, people who are told they are relatively poorer than they thought are less concerned about inequality and are not more supportive of redistribution. This finding is consistent with people using their own living standard as a “benchmark” for what they consider acceptable for others. (JEL D12, D31, H23, I31, I32)

Social commentators and researchers struggle to explain why, despite growing inequality in many countries around the world, relatively poor people are not more supportive of redistribution (Kuziemko et al. 2014, Roemer 1998, Holland 2018, Frank 2004). For example, a survey of a broadly representative sample of the national population in 44 countries by Pew Research Center shows that relatively poor respondents tend to have similar levels of support as relatively rich respondents do for raising taxes on the wealthy to fund programs to help the poor (see [Figure 1](#)) (Pew Research Center 2014). This common phenomenon would appear to contradict a key assumption underlying seminal theories of preferences for redistribution, which is that poor people should be substantially more supportive of redistribution than rich people (e.g., Meltzer and Richard 1981). Recent studies have identified

*Hoy: Australian National University (email: christopher.hoy@anu.edu.au); Mager: Oxfam Great Britain (email: fmager1@oxfam.org.uk). Dan Silverman was coeditor for this article. The study was registered with the American Economic Association RCT registry (ID numbers AEARCTR-0002534 and AEARCTR-0002614). We are grateful for comments provided by seminar participants at the Paris School of Economics, Oxford University, the London School of Economics, the University of Sydney, the Australian National University, and the Queensland University of Technology. The authors are very thankful for comments provided on an earlier version of this paper by Dean Karlan, Martin Ravallion, Michael Norton, Daniel Treisman, Edoardo Teso, Elisabeth Bublitz, Russell Toth, Eva Vivalt, Stephen Howes, Emma Samman, Mathias Sinning, Deborah Hardoon, Alice Krozer, and Ben Goldsmith. This paper presents independent analysis and is part of a project funded by Oxfam, the University of Sydney Research Excellence Initiative, and Equity Economics. The content of the paper is solely the responsibility of the authors and does not necessarily reflect the official views of the funders. The study received human ethics approval from the Australian National University (Protocol number 2017/638).

[†]Go to <https://doi.org/10.1257/pol.20190276> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

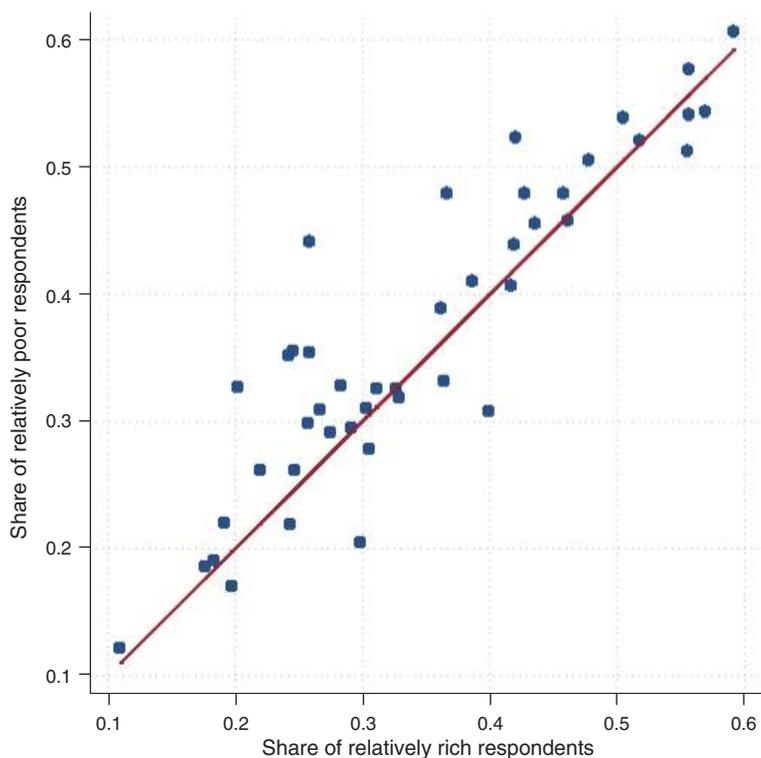


FIGURE 1. SUPPORT FOR HIGHER TAXES ON THE WEALTHY TO FUND PROGRAMS TO HELP THE POOR

Notes: This figure presents the results of a question about support for redistribution in a survey across 44 countries by Pew Research Center in 2014. Specifically, it shows the share of relatively rich and poor respondents in each country that prefer high taxes on the wealthy and corporations to fund programs to help the poor. The red line is the 45-degree line. As such, when the data point is above the 45-degree line, this means the share of relatively poor respondents that support higher taxes on the wealthy to fund programs to help the poor is larger than the share of relatively rich respondents that support the same. Relatively poor respondents have incomes below the median, and relatively rich respondents have incomes above the median. The full wording of the question is: “What would you do to reduce the gap between the rich and poor in our country? Higher taxes on the wealthy and corporations to fund programs that help the poor OR Lower taxes on the wealthy and corporations to encourage investment and economic growth?”

Source: Pew Research Center (2014)

a potential explanation for why this assumption lacks empirical support, which is that relatively poor people do not realize they are in the bottom half of the national income distribution (Gimpelson and Treisman 2018, Bublitz 2020). Surveys across a range of high-income countries have shown that most people tend to think that they are positioned around the middle of the national income distribution regardless of whether they are rich or poor (Gimpelson and Treisman 2018). This raises the question that if relatively poor people were made aware of their position in the national income distribution, would they be more concerned about inequality and supportive of redistribution?

We test how informing people that they are relatively poorer than they thought impacts their concern about inequality and support for redistribution through a randomized survey experiment with over 30,000 respondents in 10 countries

(Australia, India, Mexico, Morocco, the Netherlands, Nigeria, South Africa, Spain, the United Kingdom, and the United States). Collectively these ten countries make up around 30 percent of the global population and represent about 40 percent of world GDP. This is by far the largest survey experiment about the elasticity of people's preferences for redistribution to date and the first in multiple middle-income countries. Similar to the approach used by Alesina, Stantcheva, and Teso (2018) in a cross-country survey experiment, the sample of respondents is representative of the population with internet access in each country, and the data were collected using the survey firms YouGov, Ipsos, and Real-Time Interactive World-Wide Intelligence (RIWI). Respondents in each country were randomly allocated to receive either information about their position in the national income distribution (treatment group) or no information (control group). Prior to the treatment, respondents revealed their perception of the level of national inequality, their preferred level of national inequality, and their perceived place in the national income distribution. In addition, they provided information about their demographic characteristics, including household income that was measured using the same approach as Alesina, Stantcheva, and Teso (2018). After the treatment, respondents were asked standard questions from the existing literature regarding their views about whether the gap between the rich and poor is too large in their country and whether they think the government is responsible for closing this gap (Alesina, Stantcheva, and Teso 2018; International Social Survey Programme (ISSP) 2009). The design of our survey experiment means we are better placed to test the mechanisms through which information has an effect than the small number of previous studies were because the sample size in each country is around three times larger and we more extensively solicit people's prior beliefs.

Seminal theories of preferences for redistribution imply that informing people that they are relatively poorer than they thought would lead to greater concern about inequality and support for redistribution (e.g., Meltzer and Richard 1981, Benabou and Ok 2001, Piketty 1995). This is based on the assumption that people are averse to others in society having significantly different incomes to them and that they are more concerned about the income gap between them and the richest in society than the income gap between them and the poorest in society (e.g., Fehr and Schmidt 1999, Alesina and Giuliano 2011). However, there is little empirical support for this assumption, as only a small number of survey experiments have analyzed how information about inequality impacts people's preferences for redistribution (Hauser and Norton 2017). There is only one survey experiment that directly tests the effect of informing people that they are relatively poorer than they thought. Cruces, Perez-Truglia, and Tetaz (2013) survey 1,054 people in Buenos Aires and show that informing people that they are poorer than they thought led to greater support for redistribution, which is consistent with seminal theories. Yet, there is also some evidence from related studies that would suggest that informing people that they are relatively poorer than they thought may reduce their desire for redistribution. For example, experimental research in the United States by Kuziemko et al. (2014) shows that "last-place aversion" can exist whereby relatively poor people often prefer when there are people who are poorer than them. Similarly, there are inconsistent findings from the few studies that examine the effect of informing

people that they are richer than they thought (Nair 2018; Karadja, Mollerstrom, and Seim 2017).

We find respondents in the poorest two quintiles of the national income distribution who were told they are relatively poorer than they thought are less concerned about the gap between the rich and poor in their country and are not any more supportive of the government closing this gap compared to respondents in the control group. This result occurs in seven countries (India, Mexico, Morocco, the Netherlands, Nigeria, South Africa, and Spain), and there was no effect caused by this information in the remaining three countries (Australia, the United Kingdom, and the United States). The overall effect of the treatment was mainly driven by people who prefer low levels of inequality and by respondents in both the poorest quintile and the second-poorest quintile in most countries. We also show that there were no statistically significant effects from the treatment among respondents in the poorest two quintiles of the national income distribution who accurately estimated their position in the distribution. In addition, we conduct a series of robustness checks to illustrate that the effect was not due to how many people lived in the households of respondents, the size of respondents' misperceptions, the representativeness of the survey sample, differential attrition, or a lack of attention paid during the survey (see Section 0.1 of the online Appendix). Furthermore, we conduct an additional randomized survey experiment that replicates our original results and provides further insights as to the mechanism driving our results (see Section 0.2 of the online Appendix).

We illustrate that the likely channel causing the effect is people using their own standard of living as a "benchmark" for what they consider acceptable for others by modifying Fehr and Schmidt's (1999) seminal model of other-regarding preferences and exploring heterogeneous treatment effects. This is consistent with the results of our study as follows. People had perceived themselves to have an "average" living standard compared to other people in their country prior to the treatment even though they were actually relatively poor. Their previous assessment of their relative status implies that they thought there were similar shares of people poorer and richer than them in their country (this is as a result of placing oneself as being around the middle of the national income distribution). Upon receiving the treatment, this led people to realize two points. Firstly, there are fewer people in their country poorer than them than they had thought. Secondly, what they had considered to be an average living standard (i.e., their own standard of living) is actually relatively poor. In other words, relatively poor people consume more than they thought. Both of these points would suggest that the treatment provided to respondents would lead them to become less concerned about the living standard of poor people in their country. This notion of "benchmarking" is consistent with one of the key assumptions of most theories of preferences for redistribution, which is that relatively poor people are averse to others in society having significantly different levels of income than them. However, this study diverges from standard theory by showing that relatively poor people are more concerned about the gap between their income and that of the poorest in society than the gap between their income and that of the richest in society.

This paper contributes to the existing understanding of how people's perceptions of inequality shape their support for redistribution in at least two ways. Firstly, we identify a novel mechanism (that we refer to as "benchmarking") through which

people update their beliefs about inequality. This mechanism potentially helps to explain why poorer people are less supportive of redistribution than seminal theories of preferences for redistribution would suggest. Benchmarking means that there are competing channels through which people think about redistribution. On the one hand, relatively poor people may be more supportive of redistribution if they believe they are set to benefit, but on the other hand, they may be less supportive if they view the absolute living standard of relatively poor people as somewhat satisfactory and as such are less likely to think redistribution to the poor is needed.

Secondly, we add to the growing evidence base that suggests that seminal theories of preferences for redistribution should be modified to reflect the fact that most people do not have accurate information about the income distribution in their country (Gimpelson and Treisman 2018, Hauser and Norton 2017). We extend the stylized facts that poor people do not realize they are near the bottom of the national income distribution and that their perceived, as opposed to actual, position in the distribution is more closely aligned with their preferences for redistribution from studies in North America and Europe to a diverse group of countries (Gimpelson and Treisman 2018, Hauser and Norton 2017). However the results of our experiment illustrate that relatively poor people's misperceptions of their position in the distribution do not appear to be lowering their concern about inequality. We provide evidence that the opposite is true. Relatively poor people would be even less concerned about inequality if they knew their true position in the national income distribution. In addition, we show that the elasticity of respondents' preferences for redistribution to the treatment is substantially greater in middle-income countries than in high-income countries. This is a novel finding, as it suggests the literature on preferences for redistribution may well be qualitatively different if more research were to be conducted in these settings.

This paper is structured as follows. Section I provides a theoretical framework of how informing people that they are relatively poorer than they thought would impact their concern about inequality and explores how this framework relates to existing experimental studies on this topic. Section II outlines the methodology behind the randomized survey experiment and the econometric analysis we conduct. Section III illustrates that the descriptive trends from our survey data are consistent with previous research. Sections IV and V present the results of the randomized survey experiment and discuss how the findings relate to existing theories of preferences for redistribution.

I. Theory and Related Literature

A. Assumptions Underpinning Theories of Preferences for Redistribution

Seminal theories of preferences for redistribution tend to be based on the following three assumptions (Alesina and Giuliano 2011). Firstly, most people are averse to large income differences in their country, and this is particularly the case for relatively poor people. Secondly, people who are more concerned about inequality tend to be more supportive of redistribution. Thirdly, poorer people should be more supportive of redistribution than richer people, as they benefit directly.

To illustrate formally these assumptions about people's preferences for the distribution of income in their country, we start with a seminal model of other-regarding preferences (Fehr and Schmidt 1999).¹ For simplicity, we follow Fehr and Schmidt (1999) and assume that an individual's utility function ($U_i(x)$) is linear in inequality aversion and in their own consumption, concave and that disutility from inequality is self-centered (i.e., people care about how their consumption relates to other people's consumption but do not care about inequality between other people's consumption per se). The generalized version of this model for a set of n people can be written as follows for individual i :

$$(1) \quad U_i(x) = x_i - \beta_i \frac{1}{n-1} \left(\sum_{j \neq i} \max[x_i - x_j, 0] \right) - \gamma_i \frac{1}{n-1} \left(\sum_{k \neq i} \max[x_k - x_i, 0] \right).$$

This model captures clearly how an individual's utility depends on their own consumption (x_i) as well as the direction and size of the weighting they place on their consumption relative to people poorer (x_j) than them (β) and richer (x_k) than them (γ). Importantly, in this model, the weighting that people place on their consumption relative to others (i.e., γ and β) is independent of the size of differences in consumption. For example, in a simple three-person version of this model,² it is possible that individual i may place a much larger weight on their consumption relative to the rich person ($x_k^* - x_i$) than the poor person places on this difference in consumption ($x_k^* - x_j^*$). As such, even though the consumption gap is larger for the poor person, their level of disutility from the consumption difference between them and the rich person may be smaller than the disutility that individual i experiences (this would occur if $(x_k^* - x_j^*) / (x_k^* - x_i) < \gamma_i / \gamma_j^*$).

As mentioned above, Fehr and Schmidt (1999), along with other seminal theories of preferences for redistribution, assume that people are averse to others having significantly different consumption than them ($\gamma > 0$, $\beta > 0$) (see Alesina and Giuliano 2011). Therefore, the presence of inequality lowers people's utility even though it may not have a direct effect on their own consumption. In addition, they assume that on average, people's disutility from differences in consumption is larger for the gap between them and those richer than them than the gap between them and those poorer than them ($|\gamma| > |\beta|$). Collectively, these assumptions imply $\gamma > \beta > 0$. Therefore, in a country with a set amount of inequality on average, individuals close to the top of the distribution have higher utility (beyond just having a higher level of consumption) than those who are closer to the bottom of the distribution. As a result, for a given weighting on income differences, the poorer an individual is, the more likely they are to experience disutility from the gap between the rich and poor in their country.

Seminal theories also assume that people's concern about inequality tends to be related to their support for redistribution (e.g., Benabou and Ok 2001). As such, individuals that place a greater weighting on their consumption relative to others (i.e., they experience greater disutility from inequality) are more likely to prefer

¹ While their model refers to consumption, we follow Kuziemko et al. (2015) and focus on income in this paper.

² Where there is one person poorer than individual i and one person richer than individual i .

that the government intervenes to reduce income differences in society. However, if people have a high level of concern about inequality, this may not always mean that they have a high level of support for redistribution. For example, if an individual lacks trust in the government, they may be concerned about inequality but not support redistribution, as they believe the government will not address the problem. Therefore, changes in other-regarding preferences could be considered a necessary but not sufficient condition for preferences for redistribution to change.

Another key assumption of most seminal theories of preferences for redistribution is that poorer people should be the most supportive of redistribution, as they benefit directly. For example, the Meltzer-Richard hypothesis proposes that people below the median income in the national income distribution should support redistribution and the richest in society should not (Meltzer and Richard 1981). This is based on whether the individual is potentially set to directly benefit or lose from redistribution. In terms of the utility function above, individual i considers whether their consumption level (x_i) will be higher or lower under certain levels of redistribution, and this is what determines whether they are supportive.

B. How Correcting Misperceptions May Change Preferences

To illustrate how correcting misperceptions may change preferences, we modify Fehr and Schmidt's (1999) model to reflect recent research indicating that people misperceive their position in the distribution and that their perceptions are more closely correlated with their preferences than with what is actually the case (Gimpelson and Treisman 2018, Hauser and Norton 2017, Niehues 2014). Specifically, we show that an individual's utility is dependent on how they perceive the consumption of the other people, as opposed to the actual level of consumption of other people. The generalized model becomes

$$(2) U_i(x) = x_i - \beta_i \frac{1}{n-1} \left(\sum_{j \neq i} \max[x_i - x_j^P, 0] \right) - \gamma_i \frac{1}{n-1} \left(\sum_{k \neq i} \max[x_k^P - x_i, 0] \right),$$

whereby x_j^P is the perceived consumption of people poorer than x_i and x_k^P is the perceived consumption of people richer than x_i .

This revised model provides a framework to demonstrate how information about an individual's position in the national income distribution that corrects their existing misperceptions is likely to impact their preferences. We assume that our experiment changes the information set that respondents have available to them from I^0 to I^1 . It is important to note that we are exclusively referring to correcting people's misperceptions of their relative consumption, as it is reasonable to assume that people have complete information about their own consumption. While Fehr and Schmidt's (1999) model does not explicitly predict how preferences change when individuals are provided with information, we follow Card et al. (2012), who modify this type of utility function to illustrate how people respond to information about their relative position in a distribution. If the assumptions underlying seminal

theories of preferences for redistribution are valid, this would imply the following hypotheses:

HYPOTHESIS 1: *Informing people they are relatively poorer than they thought will increase their concern about the gap between the rich and poor in their country.*

This hypothesis is expected to hold as the information will lead people to realize the gap between their income and people richer than them is larger than they thought and seminal theory assumes people are averse to inequality and are primarily concerned about the gap between their income and people richer than them (i.e., $\gamma > \beta > 0$). This can be written formally as:

$$(3) \quad \text{if } \frac{\sum_{k \neq i} (x_k - x_i)}{\sum_{j \neq i} (x_i - x_j)} > \frac{\sum_{k \neq i} (x_k^P - x_i)}{\sum_{j \neq i} (x_i - x_j^P)} \quad \text{then } [U_i(x)|I^1] < [U_i(x)|I^0].$$

HYPOTHESIS 2: *Informing people they are relatively poorer than they thought will increase their support for the government to reduce the gap between the rich and poor.*

This hypothesis is expected to hold for two reasons. Firstly, this information will make people more concerned about inequality in their country (see Hypothesis 1) and seminal theories typically assume this will lead to a greater desire for redistribution. Secondly, this information will lead poorer people to realize they are more likely to be a direct beneficiary of redistribution from rich to poor. As such they should have a higher level of consumption (x_i) if there is greater redistribution and this will raise their utility.

If we find evidence contrary to the first hypothesis, then this would suggest that one of the key assumptions underpinning Fehr and Schmidt's model and most seminal theories of preferences for redistribution lacks empirical support. People may not be averse to others consuming significantly differently to them ($\gamma < 0, \beta < 0$) or they may experience greater disutility from the gap between their consumption and those poorer than them as opposed to the gap between their consumption and those richer than them ($|\gamma| < |\beta|$). If either of these are the case then informing people they are relatively poor will actually lower their concern about inequality because they are being told the gap between their income and the income of those below them is smaller than what they had thought (i.e., $\sum_{j \neq i} (x_i - x_j) < \sum_{j \neq i} (x_i - x_j^P)$).

It is ambiguous what the effect of informing people they are poorer than they thought will be in terms of the second hypothesis if the first hypothesis is not confirmed because this means there are competing channels through which people think about redistribution. On the one hand, poorer people may be more supportive if they are set to benefit from redistribution, but on the other hand they may be less supportive if they are less concerned about the degree of inequality in their country. The net result will depend on the size of the weighting that people place on their consumption relative to others.

C. Related Literature

Only a relatively recent and small literature of randomized survey experiments exists that analyze how people's preferences for redistribution are impacted by information (Hauser and Norton 2017). Influential studies by Kuziemko et al. (2015) and Alesina, Stantcheva, and Teso (2018) show that preferences for redistribution can be elastic to information about inequality, however neither test the hypotheses above. For example, Kuziemko et al. (2015) conduct a series of randomized survey experiments in the United States and find that only information about estate taxes was able to increase support for redistribution. The only cross-country survey experiment that relates to the research question in this paper provides a treatment with multiple pieces of information about inequality to respondents, including their position in the national income distribution (Bublitz 2020). This treatment reduced support for redistribution among people who were richer than they thought in Germany and Russia but had no effect in the remaining four countries (France, Spain, Brazil, and the United States). It is challenging to identify the mechanisms that are causing this outcome as respondents were provided with various pieces of information about inequality in their country as well as information about their position in the distribution.

There are four key studies in the literature (all of which are only based in a single country) that relate closely to the hypotheses above and their findings have direct implications about the shape of people's utility in our model. There is only one survey experiment that directly examines the impact of informing people they are poorer than they thought on their support for redistribution. Cruces, Perez-Truglia, and Tetaz (2013) present a theoretical model based on the idea that misperceptions of inequality are due to people extrapolating information from endogenous reference groups. They test this by conducting a household survey experiment of 1,054 respondents in Buenos Aires (the capital of Argentina) and show that information about people's position in the national income distribution boosts support for redistribution among people who were relatively poorer than they thought. This result (i.e., $\gamma > \beta > 0$) is consistent with the hypotheses in our paper. A related study by Kuziemko et al. (2014) provides evidence from laboratory experiments and surveys in the United States that "last place aversion" can exist whereby people near the very bottom of the distribution are the least supportive of redistribution, particularly in the form of increases in the minimum wage. It is argued that this is because poor people prefer a gap between their income and that of people relatively poorer than them. This finding (i.e., $\beta < 0$) is inconsistent with the hypotheses above.

There are two similar survey experiments that examine the effect of informing people that they are relatively richer than they thought. The first, by Karadja, Mollerstrom, and Seim (2017), draws on a postal survey experiment of 1,001 respondents in Sweden and shows that informing people they are relatively richer than they thought lowers support for redistribution³. This result is compatible with the hypotheses in our paper (i.e., $\gamma > \beta > 0$). The second, by Nair (2018),

³Almost all respondents to this study either underestimated or accurately estimated their position in the income distribution. As such, the sample size was too small (i.e., had inadequate statistical power) for the authors to examine the hypotheses above.

involves an online survey experiment of 1,559 respondents in the United States and shows that when people were told they were relatively richer than they thought in terms of global income distribution, they became more supportive of international redistribution. Nair argues that this finding is because people are averse to a large gap in incomes between them and the poorest people in the world. This finding (i.e., $\beta > \gamma > 0$) is incompatible with the hypotheses above.

II. Methodology

This study helps fill the gap in the literature about how relatively poor people's misperceptions of their position in the national income distribution impacts their support for redistribution by testing the hypotheses above.

A. Sample Selection and Sample Size

We conducted a randomized survey experiment with over 30,000 respondents in 10 countries (Australia, India, Mexico, Morocco, the Netherlands, Nigeria, South Africa, Spain, the United Kingdom, and the United States) during the last 3 months of 2017. These countries make up around 30 percent of the global population and represent about 40 percent of world GDP. The diverse set of countries was selected so as to provide confidence in the external validity of the results of the survey experiment and to ensure at least one country was included from Asia, Sub-Saharan Africa, Latin America, the Middle East and North Africa, Europe, North America, and Oceania. There is significant variation among the countries in this study in terms of level of income, inequality, government intervention in the economy, state fragility, and degree of individualism (see [Table 1](#)).

Data were collected of a representative sample of the population with internet access in each country using internationally respected online surveys firms YouGov, Ipsos, and RIWI, which is the same approach that was used by Alesina, Stantcheva, and Teso (2018). This resulted in a sample of respondents where younger people and men were overrepresented compared to a perfectly nationally representative sample, especially in developing countries (see Table OA1 in the online Appendix). Throughout the body of the paper, we present the sample average treatment effect, and to address concerns about the representativeness of the sample, we present the treatment effects weighted by the age and gender of the national population as a robustness check in the online Appendix (see Section 0.1 in the online Appendix). In general, the effects are qualitatively similar.

In every country, the treatment and control groups had at least 800 respondents, and on average, there were around 1,500 respondents in each group. This is a similar sample size in each country as to what was used in Kuziemko et al. (2015) and Alesina, Stantcheva, and Teso (2018). Previous studies that just provided information about a respondent's position in the distribution only had around 500 respondents in the treatment group (Cruces, Perez-Truglia, and Tetaz 2013; Karadja, Mollerstrom, and Seim 2017). As such, we have substantially more statistical power to detect heterogeneous effects, such as differences in the effect of information between people based upon their preexisting preferences for inequality.

TABLE 1—BACKGROUND INFORMATION ABOUT THE COUNTRIES INCLUDED IN THIS STUDY

Source	GDP per person	GINI index	Revenue	State fragility	Individualism
	(2011 PPP constant US\$)		(% of GDP)	score	index
	World Bank 2017a	World Bank 2017a	World Bank 2017a	Systematic Peace 2017	Hofstede Insights 2020
Australia	44,781	34.4	24.8	2	90
India	6,516	37.8	12.9	11	48
Mexico	17,955	48.7	20.0	5	30
Morocco	7,508	39.5	25.2	6	46
Netherlands	48,809	28.5	39.1	0	80
Nigeria	5,351	43	N/A	18	30
South Africa	12,214	63	30.9	8	65
Spain	34,125	34.7	15.7	0	51
United Kingdom	40,228	34.8	35.3	0	89
United States	54,470	41.4	19.6	3	91

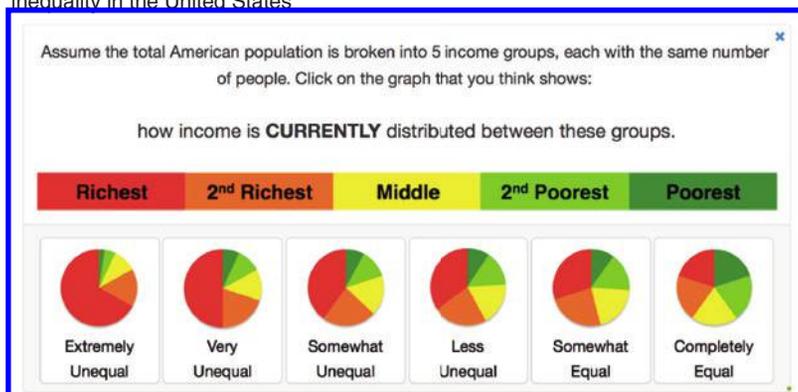
Notes: This table presents background information about the level of income, inequality, government intervention in the economy and state fragility and the degree of individualism for each of the countries in this study. According to Systemic Peace (Marshall and Elzinga-Marshall 2017), the state fragility score is a measure of both a country's effectiveness and its legitimacy in four performance dimensions—security, political, economic, and social—at the end of the year 2016 (the higher the score, the more fragile the state). It is closely associated with its state capacity to manage conflict, make and implement public policy, and deliver essential services and its systemic resilience in maintaining system coherence, cohesion, and quality of life, responding effectively to challenges and crises, and sustaining progressive development. The individualism index is a measure of the extent that one's personal identity is distinct from others', whereby a higher score means the country is more individualistic (Hofstede Insights 2020).

B. Survey Design

The survey consisted of two sections; the first collected people's existing perceptions of inequality and demographic characteristics, while the second included questions about people's concerns about inequality and desire for government action (see Section 0.3 in the online Appendix for a full list of questions). The demographic characteristics section included questions about respondents' voting preferences, their total household income, and the number of people in each respondent's household. The latter two questions were used to estimate the position of each respondent in the national income distribution. As is the case in all studies that rely on reported income, there is a risk that respondents do not provide correct information about their household's actual income (World Bank 2014). To minimize the likelihood of this occurring, we follow the example of a seminal randomized survey experiment on perceptions of inequality by Alesina, Stantcheva, and Teso (2018). Specifically, we first asked respondents how many people lived in their household. We then required respondents to select their actual household income from a list of 5 options that contained bands of income levels (e.g., less than \$10,000, \$10,000–\$20,000, etc.). We designed the survey so that these options would adjust based on the reported number of people living in each respondent's household, so that the five options corresponded with the five quintiles of the national income distribution. This type of closed question minimizes the likelihood of measurement error, as respondents were not free to enter any number, and they could review the complete set of options prior to selecting an answer.

To measure respondents' perceptions of the level of national inequality, previous studies have used a range of techniques including stylized distributions (ISSP

Panel A. Question about respondents' perception of the existing level of inequality in the United States



Panel B. Question about respondents' preference for the level of inequality in the United States

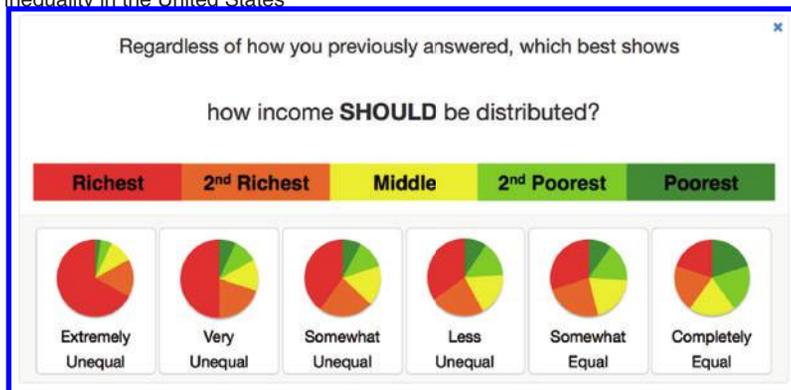


FIGURE 2. QUESTIONS ASKED ABOUT RESPONDENTS' PERCEPTION OF AND PREFERENCES FOR INEQUALITY

Notes: This figure shows the questions that were asked about respondents' perception of the existing level of inequality and preference for the level of inequality in the United States. The most unequal distribution option is based upon the actual level of income inequality in South Africa. This is followed by the distribution in the United States, the United Kingdom, and the Netherlands. The most equal distribution options are more equitable than what exists in any country in the world. The preamble to the question was adjusted in each country. For example, in the United Kingdom, the question stated "the total British population," not "the total American population." Panel A shows the question that was asked about people's views on the current level of inequality in the United States, and panel B shows the question that was asked about people's preference for the level of inequality in the United States.

2009) or asking respondents to estimate quintile shares (Norton and Ariely 2011). We follow the rationale behind existing approaches; however, we minimize measurement error by gathering people's perceptions in an ordinal sense (Kuhn 2019). Respondents were asked to select one of six options that could represent the distribution of income in their country, ranging from perfectly equal to extremely unequal (see Figure 2, panel A). They were then asked to select the level of national inequality they would prefer to exist using the same set of options (see Figure 2, panel B). The questions were identical for respondents in each country. As these questions only gather people's perceptions in a qualitative sense, we are unable to determine the extent to which people accurately estimated the level of inequality. We compare respondents' answers to the questions about their perceived and pre-

In which of these income groups do you place your household?

Richest 2nd Richest Middle 2nd Poorest Poorest

FIGURE 3. QUESTION ABOUT RESPONDENTS' PERCEIVED POSITION IN THE NATIONAL INCOME DISTRIBUTION

Notes: This figure shows the question that was asked about people's views on their perceived position in the national income distribution. This question was asked immediately after the question shown in Figure 2.

ferred level of inequality to determine whether they would prefer lower levels of inequality than what they believe currently exists.

Respondents were also asked about which quintile they perceived their household to fall into in the national income distribution (see [Figure 3](#)). This approach of using a limited number of options for respondents to select from is similar to other studies⁴ (e.g., Cruces, Perez-Truglia, and Tetaz 2013 and Karadja, Mollerstrom, and Seim 2017), as asking respondents their exact percentile or rank in the distribution is likely to have a large margin of error. As we are using the fairly coarse measure of quintiles, our estimates are likely to dramatically understate the extent to which people misperceive their position in the distribution. Therefore, the results of our survey should be considered a lower bound estimate of the effect of correcting people's misperceptions of their position in the distribution.

To determine people's misperceptions of their position in the income distribution, we compare respondents' answers to the question about their perceived quintile in the national distribution to their actual quintile in the national income distribution (based on their self-reported household income). We categorize people based on whether they accurately, under- (relatively richer than they thought), or overestimated (relatively poorer than they thought) their position. Importantly, our treatment is based on self-reported household income, so that way we can be confident that the only new component of the information we provide to respondents is their relative position. If we were to rely on another source of information for people's income, as some other studies have done (e.g., Karadja, Mollerstrom, and Seim 2017), we would not be able to be confident that the treatment effect is purely due to people being informed about their relative position. In these other studies the treatment effect could potentially be driven by the fact that the information provided from another source about people's own income is higher or lower than what they believe to be the case.

The second part of the survey included questions about respondents' concern about inequality and their support for redistribution (see [Table 2](#)). These questions were sourced from previous studies, specifically the ISSP (2009) and Alesina, Stantcheva, and Teso (2018). Importantly, the question we ask about support for redistribution is quite explicit in that it is referring to redistribution from rich to poor, which means it is very difficult to argue that respondents did not realize that the poor are meant to benefit more than the rich from redistribution (Holland 2018).

⁴The main difference is that we use quintiles in this study (as opposed to deciles) to make it easier for people to understand even if they lack basic numeracy.

TABLE 2—QUESTIONS ABOUT PEOPLE’S CONCERN ABOUT THE GAP BETWEEN THE RICH AND POOR AND WHETHER THEY BELIEVE THE GOVERNMENT IS RESPONSIBLE FOR CLOSING THE GAP

<p>GAP—To what extent do you agree with the following statement: “The gap between the rich and the poor in (COUNTRY X) is too large”? (Strongly Agree, Agree, Neither Agree or Disagree, Disagree, Strongly Disagree)</p>	<p>RESPONSIBILITY—To what extent do you agree with the following statement: “It is the responsibility of the government to reduce the gap between the rich and the poor”? (Strongly Agree, Agree, Neither Agree or Disagree, Disagree, Strongly Disagree)</p>
--	--

Notes: This table presents the questions that were asked about respondents’ concern about inequality and support for redistribution. These questions are standard in the literature and were sourced directly from existing studies. The name of the country where the survey took place was inserted where COUNTRY X is shown in the table. To analyze answers to these questions in our main regressions, we create a dummy variable that takes on the value of one if the respondent agrees or strongly agrees with the statement and takes on the value of zero otherwise.

Sources: Alesina, Stantcheva, and Teso (2018); ISSP (2009)

We also ensured that the questions we asked about people’s concern about inequality and support for redistribution were clearly linked. This allows us to easily make close comparisons between how people respond to the two questions (i.e., it is reasonable for us to compare the share of people who agree with both statements in each country), as previous studies have done (e.g., Niehues 2014).

Prior to answering the second section of the survey, respondents were randomly allocated to either receive information about which quintile in the national income distribution they belonged to (see an example in [Figure 4](#)) or no information (the control group). The treatment was designed in a way to only provide respondents with information about their relative position in the income distribution and not provide direct information about the income distribution in general. This was to ensure that we could speak directly to the hypotheses discussed in Section I and avoid complicating the interpretation of the treatment effect. Randomization ensured that the effect of information could be determined by comparing average differences in answers to questions between the treatment and control groups. There were few statistically significant differences between treatment and control groups across demographic characteristics and existing perceptions of inequality in each country (see Table OA2 in the online Appendix). Data about the income distribution in each country were sourced from the World Bank (World Bank 2017a, b) for the middle-income countries and the OECD for high-income countries (OECD 2017).

We designed our study to minimize the risk that the findings would be impacted by measurement issues that can arise in survey experiments. Firstly, there is a risk that a “placebo effect” or “priming” may occur whereby simply mentioning inequality could trigger a response irrespective of the exact content of the treatment (Nair 2018, McCall et al. 2017). We address this concern by asking all respondents prior to the treatment about their views on the level of national inequality and their perceived position in the national income distribution. As such, all respondents were already thinking about inequality (i.e., both treatment and control groups were “primed”) prior to the treatment being provided, which dramatically reduces the likelihood that there is a placebo effect from the treatment itself.

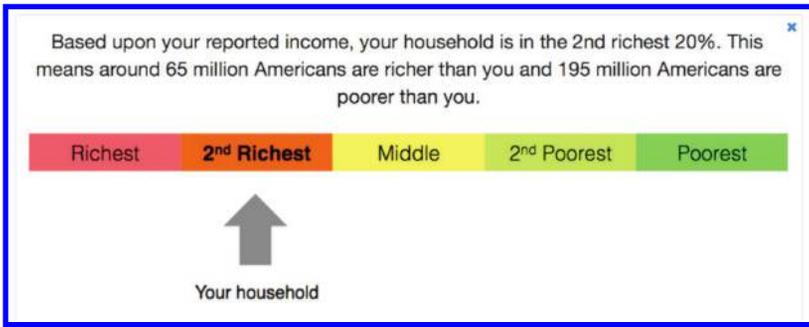


FIGURE 4. EXAMPLE OF THE INFORMATION SHOWN TO RESPONDENTS IN THE TREATMENT GROUP

Notes: This figure shows an example of the information provided to respondents in the treatment group in the United States who belonged to the second-richest quintile. The treatment was tailored to the specific circumstances of each respondent based on which country they lived in and their reported household income in per capita terms.

Secondly, a potential risk in a survey experiment is that results are skewed due to “experimenter demand effects” or “social desirability bias” (Kuziemko et al. 2015) whereby respondents provide answers they either think the experimenter would want to hear or perceive as more socially acceptable. In our study, this could look like respondents pretending to be more altruistic (measured as being concerned about inequality in their country) than they actually are. Existing studies on this topic (e.g., Kuziemko et al. 2015; Alesina, Stantcheva, and Teso 2018) have illustrated that these issues are not very prevalent in randomized survey experiments about preferences for redistribution by taking steps such as conducting follow-up surveys of a small subset of respondents a week after the original survey. They have consistently found that experimenter demand effects are not present. In our survey, we also took a number of steps to minimize the likelihood that “experimenter demand effects” or “social desirability bias” are impacting our results, as all respondents remained anonymous and the surveys were conducted online so that there was no direct human interaction associated with completing the survey. Furthermore, there was no incentive structure that could lead respondents to believe they could answer the survey in a way that increased their likelihood of getting to participate (or avoid participating) in a future survey. This is an advantage of our study over previous research that relies on platforms like Mechanical Turk where respondents know they may be reviewed based on the answers they provide to the survey (e.g., Kuziemko et al. 2015).

C. Empirical Model

We capture the effect of information by comparing average responses to the questions in Table 2 between the treatment group and the control group using an ordinary least squares (OLS) regression with a binary dependent variable (i.e., a linear probability model). This involves creating a dummy variable for each question (Y_j) in Table 2 that takes on the value one if the respondent strongly agrees or agrees with the statement in question j and the value zero if the respondent does not select

one of these options. We also create a dummy variable for the treatment group (T) that takes on the value one if the respondent belongs to the treatment group and the value zero if the respondent belongs to the control group. We preregistered the design of this study with the American Economic Association Randomized Control Trial registry (ID numbers AEARCTR-0002534 and AEARCTR-0002614) (Hoy and Mager 2017–2020).

In the body of the paper, we focus the analysis on the effect of information in each country individually because the information provided is tailored to the income distribution of each country, the questions relate to redistribution by the national government, and the political economy within each country differs. However, for completeness, we also present the main results of the survey experiment using pooled OLS regressions with country fixed effects across all countries, and we also run pooled regressions for middle-income and high-income countries separately.

Our main regression model can be written formally as follows:

$$(4) \quad Y_j = \beta_0 + \beta_1 T + X\gamma + \varepsilon,$$

where β_1 captures the average difference in the share of respondents in the treatment and control groups that agreed with the statement in question j (i.e., the treatment effect). Further, X is a vector of variables that controls for potential imbalance in background characteristics between the treatment and control groups, β_0 is the intercept term, and ε is the model error term.

We analyze the heterogeneous treatment effects from information in regards to people's actual position in the distribution, perceived level of inequality, and preferred level of national inequality. In addition, we conduct a series of robustness checks, including testing the effect of information using an ordered logit model. To do this, we follow a similar approach to the main type of analysis whereby we capture the difference in responses between treatment and control groups using the coefficient of a dummy variable for the treatment group (i.e., T).

III. Data

This section first examines whether the descriptive findings from our survey data are consistent with the assumptions underlying seminal theory, and then we examine the extent to which misperceptions about people's positions in the income distribution are present.

A. Empirical Support for Assumptions Underlying Theories of Preferences for Redistribution

This subsection illustrates that our survey data provide empirical support for most of the assumptions of seminal theories of preferences for redistribution except that they show that there is surprisingly little difference between relatively poor and rich people's desire for lower inequality, concern about inequality, and support for redistribution.

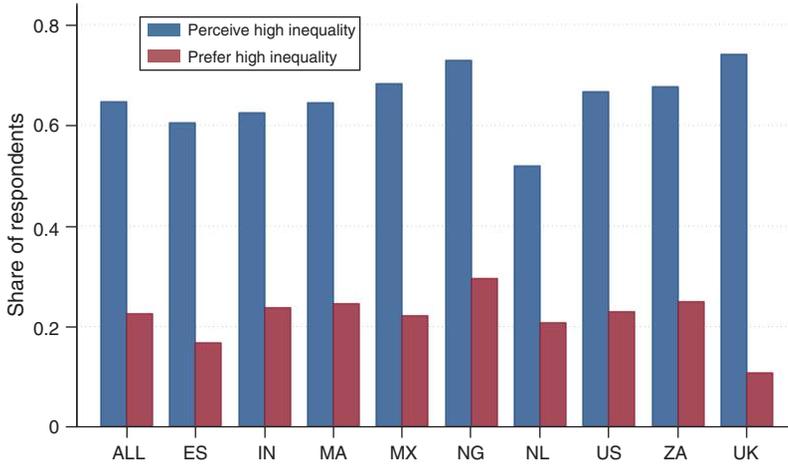


FIGURE 5. PERCEIVED AND PREFERRED LEVEL OF INEQUALITY IN EACH COUNTRY

Notes: This figure shows that more than half of people in each country perceive high levels of inequality exist but only around 20 percent of respondents prefer high inequality. ES–Spain, IN–India, MA–Morocco, MX–Mexico, NG–Nigeria, NL–Netherlands, US–United States, ZA–South Africa, UK–United Kingdom. Australia is excluded because respondents were not asked prior to the treatment about their perceived and preferred level of inequality. Prefer High Inequality–Binary variable that takes the value of one if the respondent stated prior to the treatment that they prefer high levels of inequality (defined as selecting the first or second options in Figure 3). Perceive High Inequality–Binary variable that takes the value of one if the respondent stated prior to the treatment that they perceive high levels of inequality (defined as selecting the first or second options in Figure 2).

People Are Averse to Inequality.—Consistent with seminal theories, we find compelling evidence across countries that people are averse to inequality. [Figure 5](#) shows the share of respondents in each country that perceive that high levels of inequality exist and the share of respondents that prefer high levels of inequality. Between 52 and 74 percent of respondents perceive high levels of inequality (the highest being in the United Kingdom and lowest in the Netherlands), and between 10 and 29 percent of respondents prefer high levels of inequality (the highest being in Nigeria and the lowest in the United Kingdom). However, contrary to the assumptions of seminal theories, relatively poor people tend to have similar preferences as relatively rich people for inequality to be lower than what they perceive it to be (see Figure OA1 in the online Appendix). This pattern whereby poorer people do not seem to have substantially different preferences than richer people in regards to inequality is consistent with what has been observed in other studies (Kuziemko et al. 2014, Roemer 1998, Holland 2018) and is shown in Figure 1. We find there are some countries where richer people tend to be less likely to desire lower levels of inequality (Spain, Mexico, and the Netherlands); however, there are others where the opposite is the case (India and South Africa).

Being Concerned about Inequality Is Related to Supporting Redistribution.—Across countries, there is a clear pattern whereby being concerned about inequality tends to be related to supporting redistribution (see Figure OA2 in the online

Appendix). Between 58 percent (in the Netherlands) and 79 percent (in Nigeria) of respondents agreed that the gap between the rich and poor people is too large in their country and that the government is responsible for closing this gap. An additional 11 percent (in Spain) to 23 percent (in Morocco) said that they agreed that the gap between the rich and poor is too large but did not agree that the government was responsible for closing this gap. Only a relatively small share of the population in each country (between 3 and 8 percent) stated that they support redistribution but are not concerned about inequality. These descriptive statistics are in line with the assumptions of seminal theories whereby being concerned about inequality would appear to be a necessary but not sufficient condition for supporting redistribution.

We find very little support for a key assumption in seminal theories, which is that poor people should be considerably more concerned about inequality and supportive of redistribution than rich people (see Figures OA3 and OA4 in the online Appendix). There are relatively small differences in the shares of relatively rich and relatively poor people who agree that the gap between the rich and poor is too large and that the government is responsible for closing this gap. This builds on the finding we present above whereby relatively rich and poor people tend to have similar preferences for lower inequality across countries.

B. Misperceptions of People's Position in the Distribution

This subsection illustrates that large misperceptions exist about the position of respondents in the distribution and these perceptions are correlated with people's policy preferences.

Relatively Poor People's Misperception of Their Position in the Distribution.—The vast majority of respondents from the poorest two quintiles of the national income distribution in each country were unable to accurately estimate which quintile in the distribution their household belonged to (this is calculated by drawing on responses to the question shown in Figure 3). Between 37 percent (in the United Kingdom) and 63 percent (in India) of respondents from the poorest two quintiles of the national income distribution perceived their household to be in the middle quintile of the national income distribution (see [Figure 6](#)). We call this common misperception a “median bias,” and it has also been observed in a number of studies about perceptions of inequality in high-income countries (e.g., Gimpelson and Treisman 2018). The presence of a median bias means that people who are actually poor tend to overestimate their position in the national income distribution. Between 67 percent (in the United Kingdom) and 94 percent (in Nigeria) of respondents in the poorest two quintiles overestimated their position (see Figure OA5 in the online Appendix). Only between 5 percent (in Nigeria) and 29 percent of respondents (in the United Kingdom) accurately estimated their position.⁵ Interestingly,

⁵ Only a trivial share of respondents in the second-poorest quintile underestimated which quintile in the national income distribution they belonged to. It is not possible for respondents in the poorest quintile to underestimate the quintile they belong to.

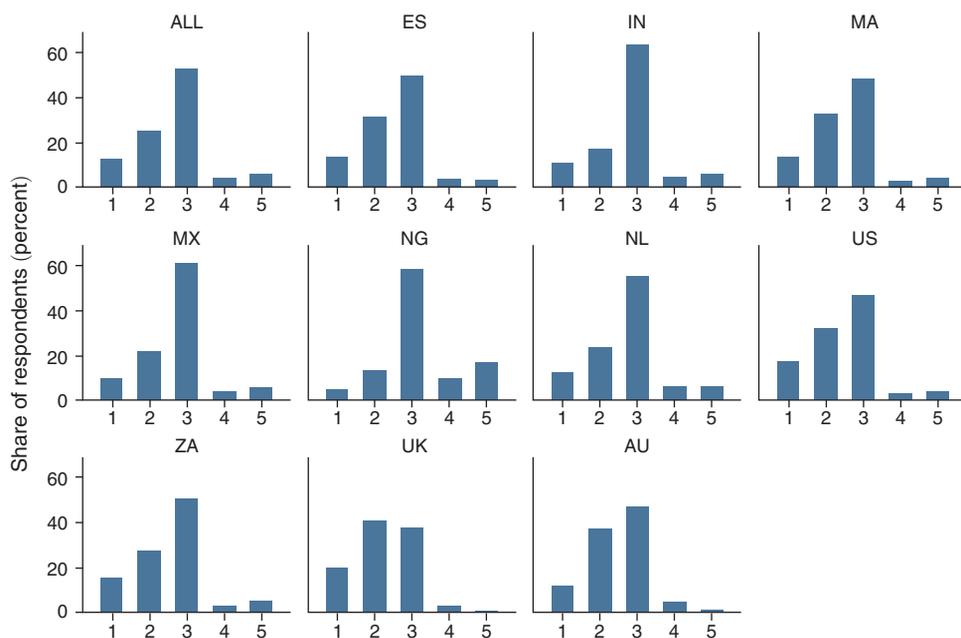


FIGURE 6. THE PERCEIVED POSITION IN THE NATIONAL INCOME DISTRIBUTION OF RESPONDENTS FROM THE POOREST TWO QUINTILES

Notes: This figure shows the perceived position in the national income distribution of respondents from the poorest two quintiles in each country, and it is based on answers to the question in Figure 3. ES–Spain, IN–India, MA–Morocco, MX–Mexico, NG–Nigeria, NL–Netherlands, US–United States, ZA–South Africa, UK–United Kingdom, AU–Australia.

respondents in high-income countries were over 50 percent more likely to accurately estimate their quintile in the national income distribution than were respondents in middle-income countries (21 percent answered accurately compared to 13 percent).

Relationship between Perceived Position in the National Income Distribution and Preferred Level of National Inequality.—The majority of respondents in the poorest two quintiles in all countries (except India) prefer lower levels of inequality than what they perceive to exist; however, this varies considerably depending on their perceived position in the national income distribution. Across all countries, relatively poor respondents (i.e., those in the poorest two quintiles) who perceived themselves to be in the poorest two quintiles of the national income distribution preferred lower levels of inequality than relatively poor people who perceived themselves to be in the middle or richest two quintiles (see [Figure 7](#)). The difference between the share of respondents in these two groups is statistically significant at a p -value of 0.05 in all countries. This finding is consistent with recent evidence in high-income countries that suggests that people’s perception of their position in the distribution is a better predictor of their redistributive preferences than what is actually the case (Gimpelson and Treisman 2018, Niehues 2014, Hauser and Norton 2017).

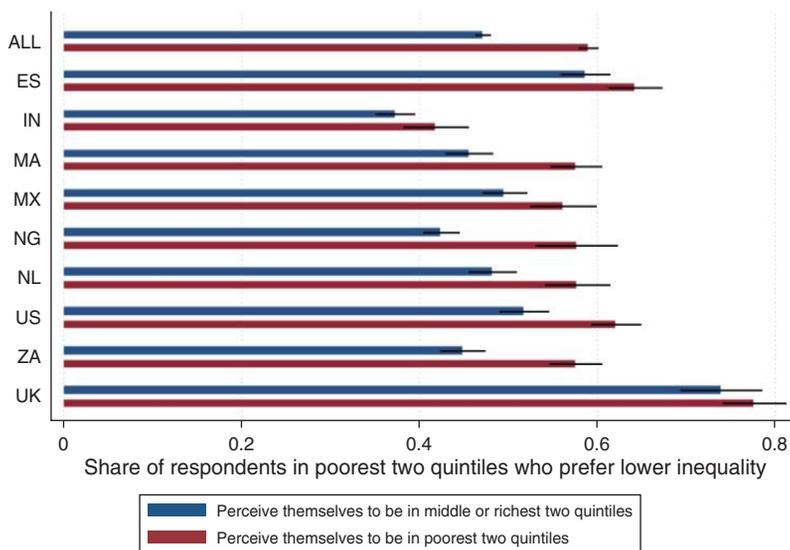


FIGURE 7. SHARE OF RESPONDENTS IN THE POOREST TWO QUINTILES THAT PREFER LOWER INEQUALITY, DISAGGREGATED BY THEIR PERCEIVED POSITION IN THE DISTRIBUTION

Notes: This figure shows the share of respondents in the poorest two quintiles who prefer lower inequality than what they perceive to exist disaggregated by whether they perceive themselves to be in the poorest two quintiles or the middle and richest two quintiles in the national income distribution. The black lines at the end of each bar present the 95 percent confidence intervals. ES—Spain, IN—India, MA—Morocco, MX—Mexico, NG—Nigeria, NL—Netherlands, US—United States, ZA—South Africa, UK—United Kingdom. Australia is excluded because respondents were not asked prior to the treatment about their preferred level of inequality. This figure shows the share of respondents in the poorest two quintiles who prefer lower inequality than what they perceive to exist disaggregated by whether they perceive themselves to be in the poorest two quintiles or the middle and richest two quintiles in the national income distribution.

IV. Results

In the previous section, we showed that relatively poor people were not substantially more concerned about inequality and supportive of redistribution than relatively rich people and presented evidence that this may be because people misperceive their relative position in the distribution. In this section, we test whether this is the case, and our results draw into question an assumption underlying seminal theories of redistribution.

A. Effect of the Treatment on Relatively Poor People Who Overestimated Their Position in the Distribution

Country-Level Effects of the Treatment.—Informing respondents in the poorest two quintiles of the national income distribution that they are relatively poorer than they thought led them to be *less* likely to agree that the gap between the rich and poor is too large in their country. This treatment effect is statistically significant in seven of the ten countries in our study (India, Mexico, Morocco, the Netherlands, Nigeria, South Africa, and Spain), and there is no effect in the remaining three countries (Australia, the United Kingdom, and the United States) (see [Figure 8](#) and [Table 3](#)).

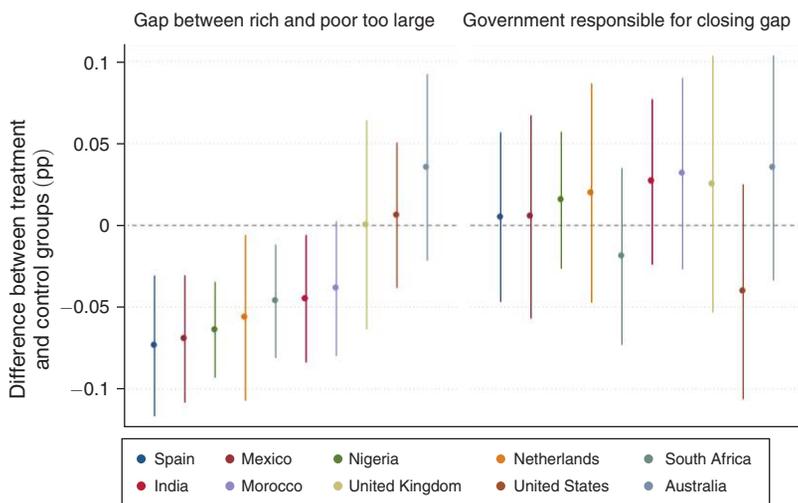


FIGURE 8. THE IMPACT OF THE TREATMENT ON RESPONDENTS IN THE POOREST TWO QUINTILES OF THE NATIONAL INCOME DISTRIBUTION THAT OVERESTIMATED THEIR POSITION

Notes: This figure provides a visual representation of the results for the main regressions in Table 3 by plotting the coefficients of the treatment dummies in each country for the outcomes of interest. The results presented in the first panel are based on the share of respondents that agree that the gap between the rich and poor is too large in their country, and the results in the second panel are based on the share of respondents that agree that the government is responsible for closing the gap between the rich and poor.

TABLE 3—THE EFFECT OF THE TREATMENT ON RESPONDENTS IN THE POOREST TWO QUINTILES OF THE NATIONAL INCOME DISTRIBUTION THAT OVERESTIMATED THEIR POSITION

	(ES)	(IN)	(MA)	(MX)	(NG)	(NL)	(US)	(ZA)	(UK)	(AU)
	b/se	b/se								
Gap too large	-0.079 (0.02)	-0.045 (0.02)	-0.037 (0.02)	-0.074 (0.02)	-0.063 (0.01)	-0.053 (0.03)	0.002 (0.02)	-0.047 (0.02)	-0.001 (0.03)	0.034 (0.03)
p-value	0.000	0.024	0.077	0.000	0.000	0.038	0.924	0.008	0.975	0.243
Mean dep. variable	0.785	0.856	0.785	0.865	0.927	0.692	0.769	0.887	0.824	0.775
Controls	Yes	Yes								
Government responsible	0.002 (0.03)	0.029 (0.03)	0.034 (0.03)	0.004 (0.03)	0.014 (0.02)	0.020 (0.03)	-0.055 (0.03)	-0.020 (0.03)	0.016 (0.04)	0.030 (0.03)
p-value	0.927	0.268	0.260	0.910	0.499	0.556	0.093	0.477	0.673	0.385
Mean dep. variable	0.776	0.807	0.661	0.717	0.823	0.647	0.654	0.768	0.682	0.609
Controls	Yes	Yes								
Observations	877	908	979	793	1,160	732	825	950	457	749

Notes: This table presents the results of the regression analysis specified in equation (4), where the dependent variable either captures the share of respondents that agree that the gap between the rich and poor is too large or the share of respondents that agree that the government is responsible for closing the gap between the rich and poor. ES—Spain, IN—India, MA—Morocco, MX—Mexico, NG—Nigeria, NL—Netherlands, US—United States, ZA—South Africa, UK—United Kingdom, AU—Australia. Gap too large—Binary variable that takes the value of one if the respondent agreed or strongly agreed that the gap between the rich and poor in their country is too large. Government responsible—Binary variable that takes the value of one if the respondent agreed or strongly agreed that the government is responsible for closing the gap between the rich and poor in their country.

Among the 7 countries where there is a statistically significant effect, the difference between treatment and control groups is between -7.9 and -3.7 percentage points. For example, in the case of Spain, 78.5 percent of respondents in the control group (who overestimated their position and are actually in the poorest 2 quintiles) agree that the gap between rich and poor is too large, whereas the share of respondents

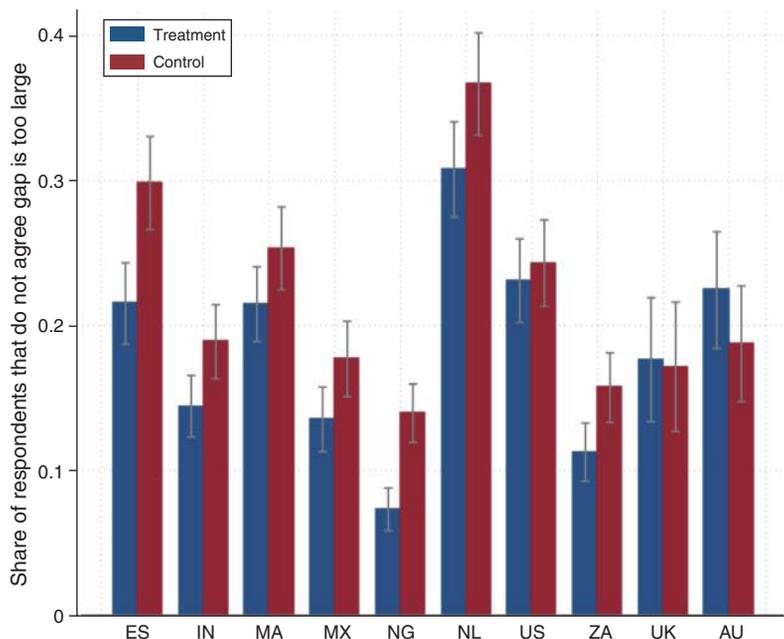


FIGURE 9. THE IMPACT OF THE TREATMENT ON THE SHARE OF RESPONDENTS THAT DO NOT AGREE THAT THE GAP BETWEEN THE RICH AND POOR IS TOO LARGE IN THEIR COUNTRY

Notes: This figure shows the share of respondents from the poorest two quintiles of the national income distribution who overestimated their position that do not agree that the gap between the rich and poor is too large in their country. ES—Spain, IN—India, MA—Morocco, MX—Mexico, NG—Nigeria, NL—Netherlands, US—United States, ZA—South Africa, UK—United Kingdom, AU—Australia.

that agree in the treatment group (who overestimated their position and are actually in the poorest 2 quintiles) is 70.6 percent. In each of the ten countries, this information does not lead to a statistically significant impact on people's views about whether the government is responsible for closing the gap between the rich and the poor.

To illustrate the magnitude of the treatment effect on people's concern about inequality, we rerun the regression analysis above with an inverted outcome variable (i.e., the share of respondents that do not agree that the gap between the rich and poor is too large in their country). We do this as concern about inequality is already very high in the control group, and this somewhat masks the impact of the treatment. We show in [Figure 9](#) that the share of respondents that do not agree that the gap between the rich and poor is too large is up to twice as high in the treatment group (in the case of Nigeria). On average, in the countries where there is a statistically significant effect, the treatment increases the share of respondents who do not agree that the gap between the rich and poor is too large by almost 40 percent.

Pooled Effects of the Treatment.—To examine the effect of the information across all countries, we also conducted pooled OLS regressions with country fixed effects. This allows us to have greater statistical power to detect the impact of the information when it exists. We show in [Table 4](#) below the results of this analysis. Consistent

TABLE 4—EFFECTS OF THE TREATMENT ON RELATIVELY POOR RESPONDENTS WHO OVERESTIMATED THEIR POSITION IN THE DISTRIBUTION POOLED ACROSS ALL COUNTRIES AS WELL AS POOLED ACROSS MIDDLE-INCOME AND HIGH-INCOME COUNTRIES

	ALL b/se (1)	MICs b/se (2)	HICs b/se (3)
Gap too large	−0.038 (0.01)	−0.045 (0.01)	−0.030 (0.02)
<i>p</i> -value	0.004	0.015	0.155
Mean dependent variable	0.859	0.890	0.822
Controls	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Government responsible	0.006 (0.01)	0.020 (0.00)	−0.011 (0.01)
<i>p</i> -value	0.484	0.002	0.449
Mean dependent variable	0.728	0.751	0.700
Controls	Yes	Yes	Yes
Country fixed effects	Yes	Yes	Yes
Observations	11,106	6,036	5,070

Notes: This table presents the results of pooled OLS regressions with country fixed effects, where the dependent variable either captures the share of respondents that agree that the gap between the rich and poor is too large or the share of respondents that agree that the government is responsible for closing the gap between the rich and poor. The first column presents the results of this analysis for all countries, the second column presents the results of this analysis for middle-income countries (India, Mexico, Morocco, Nigeria, and Mexico), and the third column presents the results of this analysis for high-income countries (Australia, Netherlands, Spain, the United Kingdom, and the United States). MICs—Middle-income countries, HICs—High-income countries. Gap too large—Binary variable that takes the value of one if the respondent agreed or strongly agreed that the gap between the rich and poor in their country is too large. Government responsible—Binary variable that takes the value of one if the respondent agreed or strongly agreed that the government is responsible for closing the gap between the rich and poor in their country.

with the country-level results, the treatment lowered concern about inequality across all countries, but primarily in middle-income countries. Interestingly, we also find a small positive effect from the treatment on support for redistribution across all middle-income countries.

B. *Heterogeneous Treatment Effects*

Main Results Disaggregated by Respondents' Prior Preferences for Inequality.—The effect of information on people's concern about the gap between the rich and poor is largely driven by respondents who stated prior to the treatment that they prefer low levels of national inequality (see Table OA3 in the online Appendix). In six of the seven countries where there is a significant effect in Table 3, the effect is driven by respondents who stated prior to the treatment that they prefer low levels of inequality (in India, Mexico, the Netherlands, Nigeria, South Africa, and Spain). Among these 6 countries, the difference between the treatment and control groups was between −8.5 and −5.9 percentage points. The exception is in Morocco, where the overall effect of the treatment is driven primarily by people who prefer high levels of inequality; among this group of respondents, the difference between

treatment and control groups was -11.7 percentage points. There are no statistically significant treatment effects in terms of respondents' views about the responsibility of the government in closing the gap between rich and poor among people who stated prior to the treatment that they prefer high or low levels of national inequality.

Main Results Disaggregated by Respondents' Prior Perceptions of Inequality.—The effect of the treatment did not vary systematically across countries by respondents' prior perceptions of the level of inequality in their country (see Table OA4 in the online Appendix). There are four countries (India, Nigeria, the Netherlands, and South Africa) where the effect is significant for respondents who perceived low levels of inequality and three countries (Spain, Mexico, and Nigeria) where the effect is significant for respondents who perceived high inequality. There are no statistically significant treatment effects in terms of respondents' views about the responsibility of the government in closing the gap between rich and poor among people who stated prior to the treatment they perceive high or low levels of national inequality.

Main Results Disaggregated by Respondents' Actual Quintile in the Distribution.—The treatment impacted respondents in both the poorest and second-poorest quintiles; however, the results are often not statistically significant (see Table OA5 in the online Appendix). The treatment effect is only statistically significant in five of the countries when restricting the sample to respondents in the poorest or second-poorest quintile. The point estimate tends to be larger in the case of respondents in the second-poorest quintile. As is the case above, there is little effect from information on people's views about the role of the government in closing the gap between the rich and poor among respondents in both the poorest and second-poorest quintiles. Nigeria is the only country where there is an effect, in which case there are opposite effects between people who are in the poorest and second-poorest quintiles.

V. Discussion

This study shows that informing people in the poorest two quintiles of the national income distribution that they are relatively poorer than they thought leads them to be less likely to agree that the gap between the rich and poor in their country is too large, and it does not have a statistically significant effect on their desire for the government to reduce this gap. These results are primarily driven by people who prefer low levels of inequality and are from both the poorest and second-poorest quintiles in the national income distribution.

A. Channel through Which Information Is Having an Effect

Our findings are clearly counter to the hypothesis that informing people that they are relatively poorer than they thought will increase their concern about the gap between the rich and poor (Hypothesis 1). We also do not find statistically significant effects in line with the hypothesis that the treatment will increase people's support for the government to reduce the gap between the rich and poor (Hypothesis 2).

To relate these findings to our model in Section I, recall that the utility function captures clearly how an individual's utility ($U_i(x)$) depends on their own consumption (x_i) as well as the direction and size of the weighting they place on their consumption relative to people poorer than them (β) and richer than them (γ). As mentioned above, seminal theories of preferences for redistribution assume people are averse to others having significantly different consumption to them and that on average, people's disutility from differences in consumption is larger for the gap between them and those richer than them than for the gap between them and those poorer than them (Alesina and Giuliano 2011). Collectively, these assumptions imply $\gamma > \beta > 0$. Our results only support the assumption that relatively poor people are averse to large differences in incomes across society. This can be seen in the descriptive statistics we show in Section III and by the fact that the treatment effect is primarily driven by relatively poor people who had expressed a prior preference for low levels of inequality (i.e., where $\gamma > 0, \beta > 0$) in all countries except for Morocco.

Our findings do not support the assumption that on average, people's disutility from differences in consumption is larger for the gap between them and those richer than them than for the gap between them and those poorer than them. When people are told that they are relatively poorer than they thought, this implies that the gap between them and people poorer than them is smaller than they thought and that the gap between them and people richer than them is larger than they thought. They respond by being less concerned about the gap between the rich and poor in their country, even though they prefer low levels of inequality. This response is only consistent with poorer people being more concerned about the gap between their income and those poorer than them compared to the gap between their income and those richer than them (i.e., $\beta > \gamma > 0$). If this was not the case, the treatment effect would be in the opposite direction.

We refer to this updating of beliefs as "benchmarking," whereby people use their own standard of living as a reference point for what they consider acceptable for others. The following explains how "benchmarking" is consistent with the results of our study. People had perceived themselves to have an "average" living standard compared to other people in their country prior to the treatment even though they are actually relatively poor. Their previous assessment of their relative status implies that they thought there were similar shares of people poorer and richer than them in their country (this is as a result of selecting oneself as being around the middle of the national income distribution). Upon receiving the treatment, this led people to realize two points. Firstly, there are fewer people in their country poorer than them than they had thought. Secondly, what they had considered to be an average living standard (i.e., their own standard of living) is actually relatively poor. In other words, relatively poor people consume more than they thought. Both of these points would suggest that the treatment provided to respondents would lead them to become less concerned about the living standard of poor people in their country. Consequently, this would result in them becoming less likely to agree that the gap between the rich and poor in their country is too large, given their utility is weighted as we show it is above (i.e., $\beta > \gamma > 0$).

To illustrate this concept of benchmarking, consider a respondent in the second-poorest quintile in India in the treatment group who, prior to the

treatment, thought they were in the middle quintile. Upon receiving the treatment, this would lead them to realize two pieces of information. Firstly, there are substantially fewer people in India poorer than them than they had thought (i.e., instead of there being over half a billion people poorer than them, they find out there are around 250 million). Secondly, what they thought was an average standard of living (i.e., their living standard) is actually relatively poor. This means that relatively poor people in India consume more than they had thought to be the case. As such, there are fewer people poorer than them than they thought, and the absolute living standard of the relatively poor is higher than they thought. Therefore, the respondent would become less concerned about the living standard of poor people in India as a result of the treatment.

B. Implications for Theories of Preferences for Redistribution

The results of our study have significant ramifications about the validity of the assumptions underlying seminal theories of preferences for redistribution. As discussed above, it is generally assumed that people who are more concerned about inequality tend to be more supportive of redistribution and that poorer people should be more supportive of redistribution than richer people, as they benefit directly. Our results suggest that these two assumptions reflect competing channels through which relatively poor people think about redistribution. This can be seen as follows. We find that informing people that they are relatively poorer than they thought lowers their concern about inequality, which seminal theories would suggest should also lower their support for redistribution. However, seminal theories would also suggest that informing people that they are relatively poorer than they thought should increase their support for redistribution, as they are set to benefit directly. Given these competing channels, it is unsurprising that we fail to detect a statistically significant effect on support for redistribution in any country.

The presence of competing channels through which relatively poor people think through preferences for redistribution implies that seminal theories need to be revised. For example, the Meltzer-Richard hypothesis (1981), put simply, proposes that people below the median income in the distribution should be the most supportive of redistribution. This is based on whether an individual is potentially set to directly benefit or lose from redistribution. However, benchmarking would imply that there is also an effect in the opposite direction. Even though relatively poor people may be likely to benefit from redistribution, they are also less concerned about the absolute living standard of people who are relatively poor and as a result are less likely to think that redistribution to the poor is needed.

This is supported by an additional survey experiment that we conducted in Australia (details in Section 0.2 in the online Appendix) that shows how there is considerable heterogeneity between respondents within countries based on their satisfaction with the status quo. Relatively poor people who were satisfied with their current income responded to the treatment by becoming less supportive of redistribution, while the opposite was the case for respondents who were unsatisfied. This is consistent with the idea that people “benchmark” their living standard relative to

others and that poor people who are more satisfied with their income are less concerned about the absolute living standard of other relatively poor people and as a result are less likely to think redistribution is needed.

The net effect on preferences for redistribution of informing people that they are relatively poorer than they thought will ultimately depend on the size of the weighting they place on their consumption relative to others. Our results suggest that on average, people place a nonzero weighting on the consumption of others relative to their own income (i.e., $\beta > \gamma > 0$); however, they are still primarily concerned about their own consumption ($1 \gg \beta, 1 \gg \gamma$). This can be seen from the fact that there is a significant negative treatment effect on concern about inequality but often a positive insignificant effect on support for redistribution in most countries. This nuanced relationship is even clearer when examining the results of the pooled regressions with country fixed effects across middle-income countries, as there is a statistically significant decrease in concern about inequality and a statistically significant increase in support for redistribution. This can be reconciled by the fact that on average, people are experiencing less disutility from inequality (the effect on concern about inequality), but on balance, they are still slightly more supportive of redistribution, as they are more likely to benefit directly.

The size of the weighting that people place on other people's consumption is likely to vary considerably across respondents and among countries. It is quite striking that the three countries that rank the highest in terms of having an individualist culture in Table 1 (Australia, the United Kingdom, and the United States) (Hofstede Insights 2020) do not have statistically significant effects from the treatment on people's concern about inequality. This is consistent with the idea that "benchmarking" is likely to be less common in countries with more individualistic cultures. In these countries, people would place smaller weights on their consumption relative to others. In addition, we find that these three countries have the highest share of relatively poor people who accurately estimate their position in the distribution in the first place (see Figure OA5 in the online Appendix). As such, there are fewer people who experience a "median bias" and subsequently fewer people for whom the treatment could correct their existing misperceptions.

C. How These Findings Relate to Previous Studies

The regularity of our findings across a diverse range of countries should be considered when assessing how they relate to previous studies. The only other study that directly tests the hypotheses in this paper, by Cruces, Perez-Truglia, and Tetaz (2013), provides evidence that when people are told that they are relatively poorer than they thought, they become more supportive of redistribution. The difference in results between their survey experiment and our study could be due to a range of factors such as differences in location or the use of a household survey as opposed to an online survey. Our study is based upon representative samples of the population with internet access in a diverse set of ten countries, whereas the sample in Cruces, Perez-Truglia, and Tetaz (2013) is restricted to people living in one city in Argentina. In addition, we use an anonymous online survey, as opposed to a household survey like Cruces, Perez-Truglia, and Tetaz (2013), to minimize the risk

of social desirability bias (when a respondent provides an interviewer with answers they think are more socially acceptable).

The notion of benchmarking that we articulate also helps to explain the findings of a recent study in the United States that shows that telling people that they are richer than they thought in the global income distribution makes them more supportive of international redistribution (Nair 2018). Nair (2018) shows that most respondents to his survey in the United States thought they only had a slightly above average living standard compared to other people around the world and that informing them that they were among some of the richest people on earth drastically changed their views. They became more supportive of foreign aid and more likely to provide money to charities operating in developing countries. This implies a similar weighting of people's utility function to our study (i.e., $\beta > \gamma > 0$), whereby people are averse to inequality and are more concerned about the gap between their income and the poorest in society than the gap between their income and the richest in society. The key difference is that we focus on the poorest quintiles in the national income distribution, while Nair (2018) focuses on the richest quintiles in the global income distribution.

Our research bolsters findings from other studies on the elasticity of people's preferences for redistribution (Kuziemko et al. 2015, Hauser and Norton 2017), which is that it is easier for information to alter people's concerns about inequality than their desire for government-led redistribution. For example, Kuziemko et al. (2015) show that providing multiple pieces of information about inequality in the United States does not lead to greater support for redistribution⁶ even though it does have a large impact on people's concern about inequality. We show a similar pattern whereby even though we are able to detect effects on people's concern about inequality from the treatment in most countries, this does not lead to statistically significant changes in their preferences for redistribution.

Our results also add considerably to the growing evidence base that suggests that seminal theories of preferences for redistribution should be modified to reflect the fact that most people do not have accurate information about the income distribution in their country (Gimpelson and Treisman 2018, Hauser and Norton 2017). Consistent with recent research in high-income countries, we show that most relatively poor people do not realize they are near the bottom of the national income distribution and that their perceived position in the distribution appears to be more closely aligned with their preferences for redistribution than their actual position in the distribution (Gimpelson and Treisman 2018, Hauser and Norton 2017, Kuhn 2019, Niehues 2014, Engelhardt and Wagener 2014). We extend these stylized facts from studies in high-income countries to a diverse group of middle-income countries. A practical implication of these findings is that policymakers who are interested in understanding people's support for redistribution in their country should be as concerned (if not more so) about people's perception of inequality as they are about what is actually the case.

⁶The exception is in regards to support for the estate tax.

VI. Conclusion

This study makes a significant contribution by showing how one of the underlying assumptions of seminal theories of preferences for redistribution lacks empirical support in a diverse range of countries. In contrast to what theories assume, informing people that they are relatively poorer than they thought reduced their concern about the gap between the rich and poor in their country. This effect was primarily driven by people who prefer low levels of inequality and actually belong to both the poorest and second-poorest quintiles in their national income distribution. We illustrate that our results are consistent with the idea that people use their own living standard as a “benchmark” for what they consider acceptable for others. This phenomenon illustrates that there are competing channels through which relatively poor people think about support for redistribution. Our findings are far more generalizable than other studies to date, as we surveyed over 30,000 people from 10 countries that make up 30 percent of the world’s population and 40 percent of global GDP.

There are three directions that we believe are promising for future research on the topic of why relatively poor people are not more supportive of redistribution. First, a deeper understanding could be developed about whether people’s preferences for redistribution are based on absolute, as opposed to relative, differences in incomes in society. Second, further survey experiments could examine what types of information, if any, would lead poorer people to become more supportive of redistribution. Finally, additional analysis could be conducted that explores the factors that shape preferences for redistribution in developing countries.

REFERENCES

- Alesina, Alberto, and Paola Giuliano. 2011. “Preferences for Redistribution.” In *Handbook of Social Economics*, edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 93–131. Amsterdam: North-Holland.
- Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso. 2018. “Intergenerational Mobility and Preferences for Redistribution.” *American Economic Review* 108 (2): 521–54.
- Benabou, Roland, and Efe A. Ok. 2001. “Social Mobility and the Demand for Redistribution: The POUM Hypothesis.” *Quarterly Journal of Economics* 116 (2): 447–87.
- Bublitz, Elisabeth. 2020. “Misperceptions of Income Distributions: Cross-Country Evidence from a Randomized Survey Experiment.” *Socio-Economic Review*. <https://doi.org/10.1093/ser/mwaa025>.
- Card, David, Alexandre Mas, Enrico Moretti, and Emmanuel Saez. 2012. “Inequality at Work: The Effect of Peer Salaries on Job Satisfaction.” *American Economic Review* 102 (6): 2981–3003.
- Cruces, Guillermo, Ricardo Perez-Truglia, and Martin Tetaz. 2013. “Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment.” *Journal of Public Economics* 98: 100–112.
- Engelhardt, Carina, and Andreas Wagener. 2014. “Biased Perceptions of Income Inequality and Redistribution.” https://diskussionspapiere.wiwi.uni-hannover.de/pdf_bib/dp-526.pdf.
- Fehr, Ernst, and Klaus M. Schmidt. 1999. “A Theory of Fairness, Competition and Cooperation.” *Quarterly Journal of Economics* 114 (3): 817–68.
- Frank, Thomas. 2004. *What’s the Matter with Kansas?* New York: Henry Holt and Co.
- Gimpelson, Vladimir, and Daniel Treisman. 2018. “Misperceiving Inequality.” *Economics and Politics* 30 (1): 27–54.
- Hauser, Oliver P., and Michael I. Norton. 2017. “(Mis)perceptions of Inequality.” *Current Opinion in Psychology* 18: 21–25.
- Hofstede Insights. 2020. “Compare Countries.” Hofstede Insights. <https://www.hofstede-insights.com/product/compare-countries/> (accessed March 17, 2020).

- Holland, Alisha.** 2018. "Diminished Expectations: Redistributive Preferences in Truncated Welfare States." *World Politics* 7 (2): 555–94.
- Hoy, Christopher, and Franziska Mager.** 2017–2020. "Cross-Country Evidence about Perceptions of Inequality and Support for Redistribution." American Economic Association Randomized Control Trial Registry. <https://doi.org/10.1257/rct.2534-13.5>. (accessed October 12, 2017).
- Hoy, Christopher, and Franziska Mager.** 2021. "Replication data for: Why Are Relatively Poor People Not More Supportive of Redistribution? Evidence from a Randomized Survey Experiment across Ten Countries." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <https://doi.org/10.38886/E128261V1>.
- Hoy, Christopher, and Russell Toth.** 2019. "A False Divide? Correcting Beliefs about Inequality Aligns Preferences for Redistribution between Right and Left-Wing Voters." Society for the Study of Economic Inequality (ECINEQ) Working Paper 494.
- International Social Survey Programme (ISSP).** 2009. "Social Inequality IV." <https://www.gesis.org/issp/modules/issp-modules-by-topic/social-inequality/2009/> (accessed July 17, 2017).
- Karadja, Mounir, Johanna Mollerstrom, and David Seim.** 2017. "Richer (and Holier) than Thou? The Effect of Relative Income Improvements on Demand for Redistribution." *Review of Economics and Statistics* 99 (2): 201–12.
- Kuhn, Andreas,** 2019. "The Subversive Nature of Inequality: Subjective Inequality Perceptions and Attitudes to Social Inequality," *European Journal of Political Economy* 59(C): 331–344.
- Kuziemko, Ilyana, Ryan W. Buell, Taly Reich, and Michael I. Norton.** 2014. "Last-Place Aversion: Evidence and Redistributive Implications." *Quarterly Journal of Economics* 129 (1): 105–49.
- Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva.** 2015. "How Elastic Are Preferences for Redistribution: Evidence from Randomized Survey Experiments." *American Economic Review* 105 (4): 1478–1508.
- Lee, David S.** 2009. "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects." *Review of Economic Studies* 76 (3): 1071–1102.
- Marshall, Monty G., and Gabrielle Elzinga-Marshall.** 2017. *Global Report 2017: Conflict, Governance and State Fragility*. Vienna, VA: Center for Systemic Peace.
- McCall, Leslie, Derek Burk, Marie Laperrière, and Jennifer A. Richeson.** 2017. "Exposure to Rising Inequality Shapes Americans' Opportunity Beliefs and Policy Support." *Proceedings of the National Academy of Sciences* 114 (36): 9593–98.
- Meltzer, Allan, and Scott F. Richard.** 1981. "A Rational Theory of the Size of Government." *Journal of Political Economy* 89 (5): 914–27.
- Nair, Gautam.** 2018. "Misperceptions of Relative Affluence and Support for International Redistribution." *Journal of Politics* 80 (3): 815–30.
- Niehus, Judith.** 2014. "Subjective Perceptions of Inequality and Redistributive Preferences: An International Comparison." https://www.relooney.com/NS3040/000_New_1519.pdf.
- Norton, Michael, and Dan Ariely.** 2011. "Building a Better America – One Wealth Quintile at a Time." *Perspectives on Psychological Science* 6 (9): 9–12.
- OECD.** 2017. "Compare Your Income." OECD. <http://www.oecd.org/statistics/compare-your-income.htm> (accessed July 17, 2017).
- Perez-Truglia, Ricardo.** 2020. "The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment." *American Economic Review* 110 (4): 1019–54.
- Pew Research Center.** 2014. "Spring 2014: Global Attitudes." Pew Research Center: Global Attitudes and Trends. <https://www.pewresearch.org/global/dataset/2014-spring-global-attitudes/> (accessed July 17, 2019).
- Piketty, Thomas.** 1995. "Social Mobility and Redistributive Politics." *Quarterly Journal of Economics* 110 (3): 551–84.
- Roemer, John.** 1998. "Why the Poor Do Not Expropriate the Rich: An Old Argument in New Garb." *Journal of Public Economics* 70 (3): 399–424.
- Schneider, Simone M.** 2019. "Why Income Inequality Is Dissatisfying—Perceptions of Social Status and the Inequality-Satisfaction Link in Europe." *European Sociological Review* 35 (3): 409–30.
- World Bank.** 2014. *A Measured Approach to Ending Poverty and Boosting Shared Prosperity*. Washington, DC: World Bank.
- World Bank.** 2017a. "World Development Indicators." Washington, DC: World Bank. <https://databank.worldbank.org/source/world-development-indicators> (accessed July 17, 2017).
- World Bank.** 2017b. "PovcalNet." Washington, DC: World Bank. <http://iresearch.worldbank.org/PovcalNet/povOnDemand.aspx> (accessed July 17, 2017).