Mind the Income Gaps? Experimental Evidence of Information's Lasting Effect on Redistributive Preferences

Bastian Becker*

September 26, 2017

Individuals reject economic inequality if they believe it to result from unequal opportunities. This paper argues income gaps between inborn groups, such as gender or race, serve people as an indication of unequal opportunities. Findings from a survey experiment show Americans underestimate these gaps. When confronted with accurate information participants correct their beliefs and adjust redistributive preferences. A follow-up survey finds these effects to last for over one year and to induce the same preference changes across the ideological divide. In sum, this paper contributes to political economy scholarship that links individual preferences to objective economic reality. Focusing on income gaps offers new ways to explore the political consequences of changes in the income distribution.

Word count: 8,976

*Bard College Berlin, Platanenstr. 24, 13156 Berlin, Germany; e-mail: b.becker@berlin.bard.edu.

I thank Thilo Bodenstein, Love Christensen, Bruno Castanho Silva, Michael Dorsch, Verena Fetscher, Achim Kemmerling, Levente Littvay, and Federico Vegetti for helpful comments and suggestions. All remaining errors are my own. Financial and institutional support was provided by Central European University (Budapest, Hungary).

Theoretical accounts linking inequality and redistribution usually assume that support for government redistribution depends on what people know about the income distribution. However, empirical inquiries into this assumption do not reveal unequivocal support for all such accounts. While it is well-established that people use information to advance their relative economic standing (Cruces et al., 2012; Karadja et al., 2016), general opposition to inequality is rare (Kuziemko et al., 2015; Trump, 2017). As most people, especially Americans, desire equal opportunities rather than equal outcomes, this is not necessarily surprising. But how does this desire relate to the income distribution? In this paper I argue that income gaps serve people as an indication of unequal opportunities, and therefore learning about them shapes redistributive preferences.

Political actors in the United States are well-aware of people's desire for equality of opportunity and appeals to it, often in terms of the "American Dream", are commonplace in political speeches and campaigns. Such appeals usually take one of two shapes. Most appeals are framed in terms of *procedural justice*, which posits that equality of opportunity is attained by subjecting everyone to the same set of rules. Critics argue that such a view does not take into account the different starting positions people are in. Instead, these critics argue to approach equality of opportunity from a point of *distributive justice*. In this view, equality of opportunity is attained only when starting positions do not affect someone's chances of success (Roemer, 1998). In American politics, the procedural perspective is commonly advanced by Republicans, whereas Democrats tend to advocate for the distributive perspective.

Political actors who favor policies that advance "distributive equality of opportunity", often promote such policies by presenting them as solutions to "gaps" between different social groups. For example, these gaps refer to systematic differences in educational and economic attainments between people of different gender, race, or migration status. The most prominent example in recent years is the White House's *equal pay* campaign under then-President Barack Obama. The campaign frequently referred to governmental statistics showing that women earn only 79% of what men receive for the same kind of work.[1] Leaving disputes over the accuracy of such numbers aside, it is not clear whether such factual information is well-suited to affect voters support of a certain policy. This is not only an important question with regards to the effectiveness of political communication. What is more, public opinion scholars frequently attest to the importance of beliefs about equality of opportunity. However, such beliefs are almost always conceptualized as being purely subjective, bearing no resemblance of objective facts, such as income gaps. To address this discrepancy–and inspired by real-world political campaigns–this paper sets out to explore how people's factual knowledge about income gaps affects their support for redistribution.

[1]See https://obamawhitehouse.archives.gov/issues/equal-pay, accessed on September 1, 2017.

This paper presents findings from an online survey experiment which asked Americans about how large they think income gaps corresponding to gender, race, and family background are. Most respondents strongly underestimate the size of the gaps. When treated with factual information those who underestimated the gaps become more likely to support redistribution, and those who overestimated them become less likely. A follow-up survey after one year shows that the effects persist. They persist equally for Republican and Democratic respondents, indicating that concerns about distributive equality of opportunity are shared across the ideological divide. Interestingly, the factual information has no lasting effect on Democrats' knowledge about income gaps, whereas treated Republicans express more accurate views in the follow-up survey. In sum, this paper develops and provides evidence for a new mechanism about how objective characteristics of the income distribution affect redistributive preferences. The paper also shows that–contra to frequent criticism of survey experiments–informational effects can be long-lasting.

Inequality, Information, and Redistribution

Political economy models posit that people form preferences for government *re*distribution based on a given distribution of economic resources. It is most commonly argued people support redistribution if they benefit from it economically. This assumption of self-interest also underlies the canonical model in the literature (Romer, 1975; Meltzer and Richard, 1981). In that model, individuals are more supportive of redistribution if they directly benefit from tax transfers, i.e. low-income earner support redistribution and high-income earners oppose it. At the same time, low as well as high-income earners become more supportive of redistribution if inequality is high. This mechanism is rooted in the effect of taxation on work incentives. This effect is most detrimental in low inequality environments, which can lead even low-income earners, who benefit directly from transfers, to oppose redistribution.

Several empirical studies explore the relationship between inequality and individuals' support for redistribution. While Finseraas (2009) finds evidence for the hypothesized positive relationship among European countries, cross-sectional studies with a broader scope of countries find little or no evidence of any relationship (Lübker, 2007; Dion and Birchfield, 2010; Haggard et al., 2013). However, cross-sectional studies are often seen with skepticism. In such non-random sets of geographic units, many contextual factors other than inequality vary and risk confounding estimation results (Milanovic, 2000). Therefore, panel studies that observe geographic units at multiple points in time, and thus are able to control for such contextual factors, are more adequate. Following this approach, studies of European countries (Schmidt-Catran, 2016) as well as European regions (Rueda et al., 2014) provide further evidence of a positive relationship between inequality and support for redistribution. Similarly, analyzing changes in inequality across US states between 1978 and 2010, Dimick et al. (2017) also present results that indicate a positive effect on support for redistribution. While not all empirical inquiries attest to a relationship between inequality and individuals' support for redistribution, those with stronger research designs do. As such, the empirical literature affirms the positive relationship Romer (1975) and Meltzer and Richard (1981) hypothesized.

However, not all scholars agree that it is concerns about work incentives that drive the relationship between inequality and support for redistribution. For example, Rueda et al. (2014) suggest that it is fear of crime that leads high income earners to become more supportive of redistribution as inequality grows. That being said, the most prominent alternative explanation, inequity-aversion, was popularized by Fehr and Schmidt (1999). Inequity-aversion suggests people reject unequal outcomes if the causes of these outcomes are illegitimate. Fehr and Schmidt (1999) argue that subjects in experiments enter as equals, because physical separation makes it impossible for subjects to identify any differences. Therefore, any deviation from an egalitarian outcome (i.e. same outcome for all) must be illegitimate. As a result, in experimental settings inequity-aversion reduces to inequality-aversion. The break-through of Fehr and Schmidt (1999) was to show that behavior in economic experiments, such as the prisoners' dilemma or the dictator game, can often be better explained with inequality-aversion than with self-interest.

Other scholars furthered this line of research with experiments specifically tailored to the study of political behavior and attitudes. Sauermann and Kaiser (2010) set up an experiment where a group of participants gets to decide over the pay-off structure through iterated voting. They show that a group's decision can be better explained if in addition to self-interest, one considers that individuals also seek to minimize the difference between their income and the group's mean income, an indication of inequality-aversion. As mentioned above, that people are inequality-averse in experimental setting hinges on the assumption that they enter the experiment as equals and that the experiment itself provides no way to legitimize outcome differences. Two studies have shown that when such ways exist, subjects become less averse to outcomes differences (Krawczyk, 2010; Balafoutas et al., 2013). In the respective experiments, subjects are given the opportunity to redistribute pay-offs just before the end of the experiment. Subjects are less likely to do so if unequal outcomes are allocated based on performance in a task and more likely to redistribute if outcomes have been allocated through a lottery.

What do these experimental insights tell us about the outside world? This outside world is characterized by large and growing economic inequality, especially in the United States (Lindert and Williamson, 2016). Unlike in the experimental setting, inequity-aversion does not reduce to inequality-aversion. To understand inequity-aversion in the outside world, one has to take into account what people know about the factors contributing to unequal outcomes, and which of these factors they regard legitimate and which not. It is widely acknowledged that individuals reject outcome differences due to factors beyond individual control and accept differences resulting from factors within individual control. In other words, people are seen to reject inequality in opportunity. For example, McCall (2013) conducts an elaborate analysis of American public opinion over the past 30 years, a time during which inequality increased strongly.[2] She argues that Americans growing opposition to inequality is not a concern about inequality itself but rather about narrowing opportunities. This echoes earlier findings by Fong (2001), who explores how beliefs about the determinants of poverty and wealth influence preferences for redistribution. She finds that those who believe in the importance of effort oppose redistribution, and those who believe in the importance of circumstances and luck support it (see also, Linos and West, 2003).

A few recent studies set out to connect individuals' concern about unequal opportunities to objective characteristics of the income distribution. These studies move beyond the focus on non-verifiable subjective beliefs. They focus on intergenerational mobility, which describes how strongly socio-economic standing of parents and their offspring coincide. Jaime-Castillo and Marqués-Perales (2014) conduct a survey among Spanish respondents and find that those who believe mobility to be lower are more supportive of redistribution. This finding is corroborated by a study which explores the beliefs about economic mobility in five countries (Alesina et al., 2017).[3] That being said, what these studies also show is that people generally hold inaccurate beliefs. Whereas Americans tend to overestimate intergenerational mobility, Europeans are prone to underestimating it.

This is no surprise. Biases in individual beliefs about objective characteristics of the income distribution are well established. Studies across a range of countries show that most people, including the poor and the rich, think of themselves as middle class and that their incomes are close to the national average (Evans and Kelley, 2004; Cruces et al., 2012; Fernández-Albertos and Kuo, 2015; Kuziemko et al., 2015). Similarly, people tend to underestimate the overall extent of inequality (Osberg and Smeeding, 2006; Norton and Ariely, 2011) or simply hold inaccurate beliefs (Loveless and Whitefield, 2011).

One explanation for deviations of beliefs from objective characteristics is limited information. A number of recent experimental studies show that inaccurate beliefs, and corresponding preferences, can be corrected through the provision of factual information. For example, if people are exposed to their actual income rank, those who previously overestimated it become more supportive of redistribution (Cruces et al., 2012; Karadja et al., 2016), especially through progressive taxation (Fernández-Albertos and Kuo, 2015).[4]

[2]See also McCall and Kenworthy (2009).

[3]These countries are France, Italy, Sweden, the United Kingdom, and the United States.

[4]The first study is based on a sample of Argentinian respondents, the latter on a Spanish sample.

Those who underestimated it, and thus realize that they are relatively richer than they thought, become less supportive of such redistributive policies. Kuziemko et al. (2015) confront American respondents with an 'omnibus treatment' that contains information about the extent of inequality and its recent growth. They find that the treated upwardly adjust their beliefs about inequality and increase their support of government redistribution.

Information also plays an important role when it comes to beliefs about intergenerational mobility. In fact, the study by Alesina et al. (2017) features a large-scale, comparative survey experiment. Providing factual information about intergenerational mobility leads to a correction of belief biases and, at least among left-leaning individuals, to greater support of redistribution. In sum, this and other studies underline that information condition beliefs about objective characteristics of the income distribution and corresponding preferences.

Contribution

As elaborated above, it is well-established in the literature that individuals reject outcome differences if they believe them to be the result of factors *beyond* individual control. However, it is largely unclear how such beliefs relate to the actual income distribution. One exception is research on intergenerational mobility. This research has shown that objective information about chances of upward mobility is an important determinant of respective beliefs (Alesina et al., 2017). Of course, parents' economic standing is of central importance for offspring's economic opportunities, but there are many other consequential factors beyond individual control. This applies to many characteristics determined at birth, like gender or race. Since such characteristics are invariably unaffected by individual choices, they constitute factors beyond individual control. As for intergenerational mobility, income differences corresponding to inborn characteristics, or income gaps for short, indicate the presence of unequal opportunities.

It is important to consider that income gaps are an imprecise indicator for unequal opportunities. In fact, income gaps might be the result of compositional differences in factors *within* individual control, such as effort. That being said, income gaps do convey some information about the possibility of compositional differences. To demonstrate this point, consider that there is only one factor beyond individual control, B, and several other factors within individual control, one of them being W. Both factors, B and W, are binaries taking either value 0 or 1. If income is indicated by Y, $\mathrm{E}(Y|B=1) - \mathrm{E}(Y|B=0)$ constitutes the overall income gap. However, this income gap is an imprecise indicator of unequal opportunities since the gap might be due to compositional differences of W, or other factors within individual control, across levels of B. This imprecision is reduced if income gaps are determined for each level of W separately, i.e. $\mathrm{E}(Y|B=1 \wedge W=1) - \mathrm{E}(Y|B=0 \wedge W=1)$

and $\mathrm{E}(Y|B = 1 \wedge W = 0) - \mathrm{E}(Y|B = 0 \wedge W = 0)$. These gaps can still be due to compositional differences in factors other than W, but not W. This example demonstrates that the uncertainty inherent to income gaps as an indicator of unequal opportunities decreases as more factors within individual control are accounted for (see also Oaxaca, 1973).

What do people know about income gaps and how does it influence their redistributive preferences? Prior research shows that people tend to underestimate income differences, i.e. the distance of their income from the national average and the extent of inequality in general, and overestimate the extent of intergenerational mobility (at least Americans). I expect the same to hold for income gaps, *people tend to underestimate income gaps.* Furthermore, as income gaps are an indication of unequal opportunities, I expect that *people who believe in larger income gaps to be more likely to support redistribution.* Finally, if beliefs about income gaps are constrained by the available information, new information should lead to a correction of belief biases and an adjustment of redistributive preferences. Concerning beliefs about income gaps I propose that *information about income gaps leads people who underestimated them to correct their beliefs upwardly and those who overestimated them downwardly.* In particular, I contend that people incorporate the new information through Bayesian updating, which implies that posterior beliefs constitute a weighted average of prior beliefs and the new information (Griffiths et al., 2008). With regards to redistributive preferences I propose that *information about income gaps makes people who underestimated them more likely to support redistribution and those who overestimated them less likely.*

The Role of Ideology. It is sometimes argued that when it comes to the formation of political preferences economic concerns are overwritten by ideological ones (Evans and Andersen, 2006). It is well-known that Americans have widely different views on the proper role and size of government. Reflecting differences in ideological predispostions, these views diverge sharply along partisan lines. While Democrats prefer a larger and more active role of government, Republicans are of the opposite opinion. Thus, Democrats are commonly supportive of increased governmental redistribution and Republicans are not (Bartels, 2008). What is more, a changing economic environment or people's knowledge about it might do little to change views on ideologically contentious topics such as governmental redistribution. That being said, the above-discussed findings from published survey experiments largely speaks against this. Experimental subjects adjust their redistributive preferences when confronted with information about their economic standing, inequality, or intergenerational mobility.[5] However, it is not clear from these experiments whether

[5]That being said, many areas in the social sciences are currently facing a "replication crisis" (Collaboration, 2015; Camerer et al., 2016). It is increasingly acknowledged that the frequent inability to reproduce established experimental findings is the result of a publication bias against null findings. Until

information leads to a lasting preference change or people quickly return to ideologically more consistent positions (see also, Gaines et al., 2007).

Furthermore, informational effects can be mediated by ideological predispositions. For example, Karadja et al. (2016) find that it is conservative Swedes who react to new information about their economic rank. Those conservatives who find out they are richer than they thought reduce their support for redistribution. New information has no effect on the preferences of liberal Swedes. Sides (2016) explore the relationship between Americans' knowledge about the estate tax, which is a tax paid on inherited property, and their support for it. Most Americans do not realize that less than 2% of the general population are directly affected by the tax. When informed about this, Republicans become more supportive of the tax whereas Democrats do not change their support. However, not all experiments indicate that Democrats, or those with liberal views, are immune to factual information. In their comparative experimental study about intergenerational mobility, Alesina et al. (2017) show that both conservatives and liberals adjust their beliefs about mobility when treated with information. That being said, effects on redistributive preferences are limited to liberals, and the authors speculate that they might simply be "preaching to the choir". As such, liberals/Democrats might be more responsive to information about economic conditions that concern the general public, whereas conservatives/Republicans might be more responsive to information about their own economic well-being.

Taking income gaps as an indication of unequal opportunities implies a distributive understanding of equality of opportunity. While Democrats in the US tend to subscribe to this perspective, Republicans are more likely to invoke a procedural understanding of equality of opportunity. The procedural understanding is not concerned with income gaps as it is the equal application of rules that is decisive, not the effect of rules on outcomes. Due to this discrepancy I expect that *Republicans and Democrats equally update their beliefs about income gaps when facing new information*, but I do not expect this to be the case for redistributive preferences. Since Republicans do not consider outcomes relevant for assessing equality of opportunity, I expect that *Republicans' redistributive preferences are not affected by their knowledge of or new information about income gaps.* The same cannot be said for Democrats. Because they subscribe to distributive equality of opportunity, I expect that *Democrats' redistributive preferences are affected by their knowledge of or new information about income gaps.*

The next section lays out how these propositions are to be tested using a survey experiment that treats respondents with information about income gaps that reflect how labor market returns differ by gender, race, and parental education.

this bias is remedied, consumers of scholarly literature have to be careful to not interpret the absence of published null finding as support for a hypothesis that has found support in only a few experiments.

The Income Gaps Experiment

Survey experiments that explore the causal effects of factual information have grown increasing popular in the social sciences. For example, Kuklinski et al. (2000) investigate people's knowledge about the amount of welfare expenditure and Hopkins et al. (2016) knowledge about the size of the immigrant population. However, most such survey experiments do focus on people's knowledge about the income distribution (Cruces et al., 2012; Fernández-Albertos and Kuo, 2015; Kuziemko et al., 2015; Trump, 2017).

While such experiments can be incorporated into face-to-face or telephone surveys, they are increasingly conducted over the internet, in particular the online time sharing platform MTurk (e.g. Kuziemko et al., 2015; Hainmueller and Hopkins, 2015; Trump, 2017). Through this platform "requesters" can offer tasks for pay to a pool of registered "workers". Academics use this platform to recruit participants for their online studies. The advantages of such online studies are not only speed and affordability, but online platforms often provide a broader pool of respondents than the more commonly used student samples. Furthermore, numerous evaluation studies of MTurk show that established findings of experimental studies and economic games can be reliably replicated (Berinsky et al., 2012; Mullinix et al., 2015; Clifford et al., 2015).

An important challenge for survey experiments that use information treatment is the interpretation of revealed effects. While a properly implemented experiment provides evidence for the absence or presence of a treatment effect, it is not necessarily the factual information itself that underlies the effect. One concern is social desirability, which describes how participants adjust their behavior and responses to what they think is expected and appropriate (McDermott, 2002). This risk is high for research about contentious topics, such as inequality. Another concern is priming. Rather than considering its factual content, an informational treatment can lead people to think about subsequent choices and answer in a particular way. For example, information about income differences primes economic concerns rather than ideological ones and this affects people's subsequently stated preferences for redistribution (Kuklinski et al., 2000). As such, priming and social desirability potentially confound any revealed treatment effects and thus threaten the internal validity of survey experiments interested in effects of factual information.

These threats to internal validity can be overcome by adjusting the research design and analysis. As factual information about income differences cannot avoid priming economic considerations, it is important to equally prime those in the control group who receive no factual information. Most survey experiments on inequality do this by asking all participants about their knowledge of economic facts (i.e. their own standing or inequality in general). Correct information is then only provided to those in the treatment group. Since asking all participants about their prior knowledge gives all of them an idea of what the survey is about, doing so has the additional advantage of minimizing social desirability biases. Further precautions can be taken during the analysis. To understand how this is done, it is important to consider that treatment effects should depend on participants' prior beliefs. Those in the treatment group who learn most from the factual information should adjust their beliefs and preferences most strongly. Hence, the presence of such an interaction effect in the analysis is a strong indication that it is the factual content of the treatment that explains its effect (Kuklinski et al., 2000; Lenz, 2009). As most above-mentioned experimental studies of economic inequality, the present paper follows these best practices to avoid confounding through priming or social desirability biases.

In specifying the income gap treatment, this study focuses on three social divides that are frequently subject to academic and public debates, gender, race, and family background. The latter divide is akin to intergenerational mobility, which has been the subject of earlier survey experiments, and I thus refer to is as intergenerational gap. Here, the intergenerational gap distinguishes the incomes of those who have at least one university-educated parent and those who do not. The race gap indicates income differences between whites and non-whites.

As discussed above, such income gaps provide an imprecise signal for the presence of unequal opportunities. This is because gaps might be the result of compositional differences in factors within individual controls. This imprecision equally applies to earlier studies on intergenerational mobility (e.g. Jaime-Castillo and Marqués-Perales, 2014; Alesina et al., 2017). In this paper, to increase the precision of the signal, I focus on income gaps only among individuals that are currently employed. These income gaps give an indication of unequal opportunities in the labor market and thus cover the larger part of the adult population. At the same time, excluding incomes of those who are currently unemployed reduces the impact of factors within individual control, such as lack of effort, skills, or choices to abstain from the labor market.[6]

The size of these gaps has been calculated for the year 2010 based on data from the Panel Study of Income Dynamics. In order to reflect labor market differences and not the redistributive effects of taxation, before-tax income data was used. This data includes income from both employment and self-employment, but not income from property or other investments. All incomes were adjusted for life-cycle variations by correcting for systematic differences based on a cubic regression of income on age (see Becker, 2017). The resulting gender gap amount to US$27,300, the race gap to US$17,800, and the intergenerational gap to US$18,700.

As discussed above, prior knowledge about income gaps in the labor market is elicited

[6]Future studies might want to use more precise signals by accounting for factors such as education, occupation, or working hours. However, a less precise signal was chosen here in order to keep the presentation of the informational treatment as simple as possible.

Figure 1: Interface to Elicit Prior Knowledge of Income Gaps.

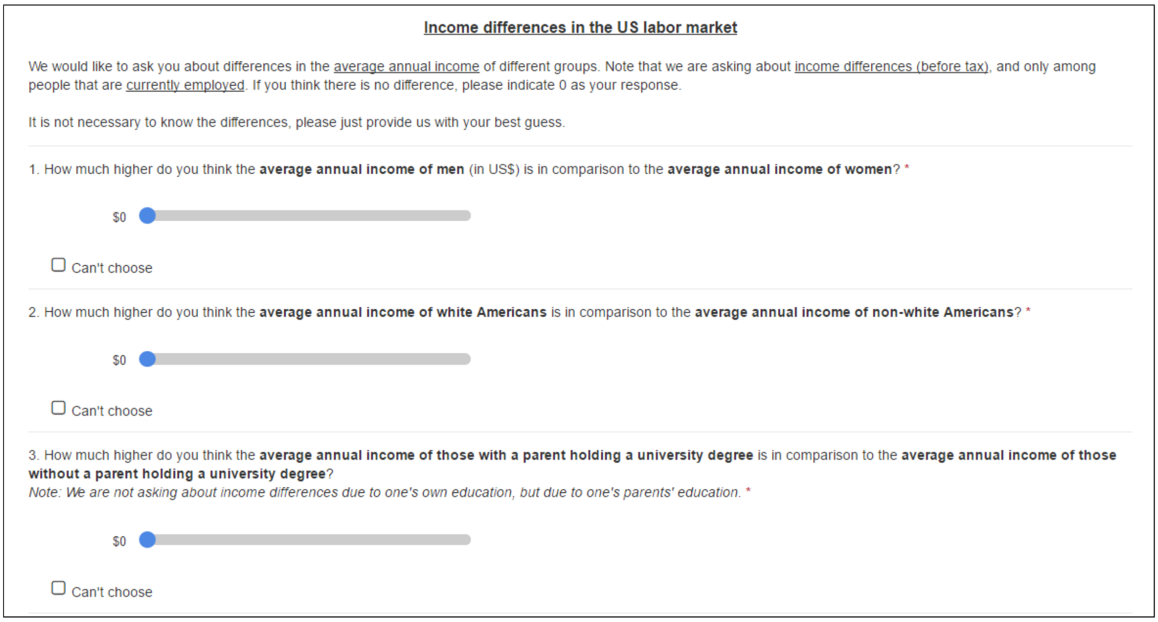

Note: Respondents can indicate any multiple of 250 between US$0 and 37,500.

for respondents in treatment and control group. Figure 1 shows the interface which is used for this purpose; respondents can drag the slider to any multiple of 250 between US$0 and 37,500. Once respondents indicate their guesses, those in the control group immediately proceed to the post-treatment questions, whereas those in the treatment group are presented with the factual information before proceeding. They are presented with the information in the same interface by additional dots on the sliders. These dots are red if the respondent underestimated the income gap and green if she overestimated the respective gap. This is complemented by a short text above each slider stating whether the respondent's guess was below or above the actual value.

Another point of contention for survey experiments is the duration of treatment effects. As follow-up surveys are rare, scholars are skeptical that effects last (Gaines et al., 2007). Kuziemko et al. (2015) constitutes one such exception. One month after their experiment, which included an 'omnibus treatment' with numerous facts about income and wealth disparities, they conduct a follow-up survey. Encouragingly, they still discover statistically significant differences between both groups for most variables of interest, including support of governmental redistribution. However, one concern about their findings is the low response rate of the second survey; only 14% of the original respondents participated. If response patterns differ between experimental conditions, so-called attrition bias, comparisons of control and treatment group cannot be interpreted causally anymore. Kuziemko et al. (2015) identify such attrition bias in their sample and are careful in drawing strong conclusions.

To explore whether information about income gaps has a lasting effect on redistributive

Table 1: Descriptive Statistics of Respondent Sample (Survey 1).

	Mean	St.Dev.	Min.	Max.	Miss.(N)
Male	0.54	0.50	0.00	1.00	0.00
Income	3.72	3.59	0.05	19.50	0.00
Age	36.39	11.68	19.00	71.00	0.00
Household size	2.67	1.49	0.00	10.00	0.00
Children	0.78	1.17	0.00	5.00	1.00
University educated parent	0.58	0.49	0.00	1.00	0.00
Race: White	0.80	0.40	0.00	1.00	0.00
Black	0.06	0.24	0.00	1.00	0.00
Other	0.14	0.35	0.00	1.00	0.00
Employment status: Unemployed	0.08	0.28	0.00	1.00	0.00
Full-time	0.64	0.48	0.00	1.00	0.00
Part-time	0.13	0.34	0.00	1.00	0.00
Keeping house	0.04	0.20	0.00	1.00	0.00
Retired	0.04	0.19	0.00	1.00	0.00
Student	0.05	0.21	0.00	1.00	0.00
Other	0.02	0.15	0.00	1.00	0.00
Education: Less than high school	0.01	0.07	0.00	1.00	0.00
High school	0.36	0.48	0.00	1.00	0.00
University	0.63	0.48	0.00	1.00	0.00

Note: Income in US$10,000. University educated parent is a dummy variable indicating whether respondent has at least one parent with a university degree. Education refers to respondent's education.

preferences, I conducted a follow-up survey with respondents after one year. In order to increase the response rate, respondents who volunteered their e-mail address in the initial survey were invited to a paid follow-up survey. This strategy proofed successful, leading to a high response rate and no detectable attrition bias (details below). Following questions about redistributive preferences, the second survey also asks respondent to again guess the size of income gaps. The results section shows that the treatment does indeed have a lasting effect on beliefs about income gaps as well as preferences for redistribution.

Re-surveying respondents after one year has further advantages. Most workers on MTurk complete academic surveys on a frequent basis (Stewart et al., 2015). In contacting those who provided their contact information, any explicit reference to the initial survey was avoided. Recipients are only informed that they are being contacted because they "previously participat[ed] in one of our surveys." The only information recipients could use to connect the message to the earlier survey is the e-mail address through which they are contacted. This seems very unlikely. As a result it is equally unlikely that any priming effects or social desirability biases induced by the income gap treatment are still at work during the follow-up.

The initial survey which included the treatment and asked about preferences for redistribution was fielded in two rounds. The first in May 2016, the second in June 2016. The survey received a total of 441 responses. Due to duplicate IP addresses, failed attention checks, or lack of permanent residence in the US, the analysis is restricted to 371 of them.[7] Randomization led to 194 of these respondents being in the treatment group and 177 in control. While the pool of MTurk workers cover a wide range of socio-demographics, samples drawn from it are not representative of the US population. Table 1 shows the composition of the present sample. Similar to related studies participants are disproportionately white, young, university-educated, non-religious, and have fewer children than the average American. More importantly though, these covariates are well-balanced across both conditions (not shown here).

Respondents who provided their e-mail address during the initial survey were contacted in August 2017. They were offered to participate in the follow-up survey for pay. In case they had quit the MTurk platform, the e-mail invited them to complete the follow-up without pay. Of the 317 respondents who earlier provided their e-mail address, 117 filled in the paid survey and 30 participated in the unpaid one. This equals a response rate of 46.4%,[8] implying that 39.6% of the initial respondents were successfully recruited for the follow-up.

Results

In line with prior studies that find Americans to underestimate the extent of economic disparities when it comes to their own income and inequality in general, I find that most participants recruited for this study vastly underestimate income gaps in the labor market (see Figure 2). A total of 98.1% of respondents underestimate the gender gap, 80.1% the race gap, and 69.6% the intergenerational gap. Participants are least off on the intergenerational gap, that is income differences due to parents' education, where the mean guess amounts to US$14,142. The mean guess of the race gap is US$11,287 and US$9,037 for the gender gap. In the following, the average guess of all three gap serves as indicator of each respondent's prior knowledge about income gaps. The distribution of prior knowledge about income gaps has a mean of US$11,459 (St.Dev.: 6,323).

To analyze the effects of the informational treatment, I dichotomize the main dependent variable, preferences for redistribution, to distinguish those who agree with increased governmental redistribution from those who do not.[9] Doing so avoids having to make the assumption that respondents interpret all seven answer categories of the survey question, ranging from strongly agree to strongly disagree, in the same way.[10] Furthermore,

[7] I also excluded three respondents who guessed none or only one income gap and another two who did not respond to the redistribution question.

[8] This is the percentage of those who provided their e-mail in the first survey. A small amount of messages was returned due to incorrect or expired addresses,

[9] *Neither agree nor disagree* responses are subsumed under the latter category.

[10] The results presented in the following are not driven by this specification. Regression tables in

dichotomous dependent variables in experiments can be analyzed with linear probability models, whose functional form fits the theoretical conjecture that individuals are Bayesian updaters. Alternative link function which are often used for dichotomous dependent variables, such as logit or probit, would necessitate different assumptions about how individuals incorporate new information. All in all, the use of a linear probability model in combination with a dichotomous dependent variable implies that the estimated model

Figure 2: Prior Knowledge of Income Gaps before Treatment (Initial Survey).

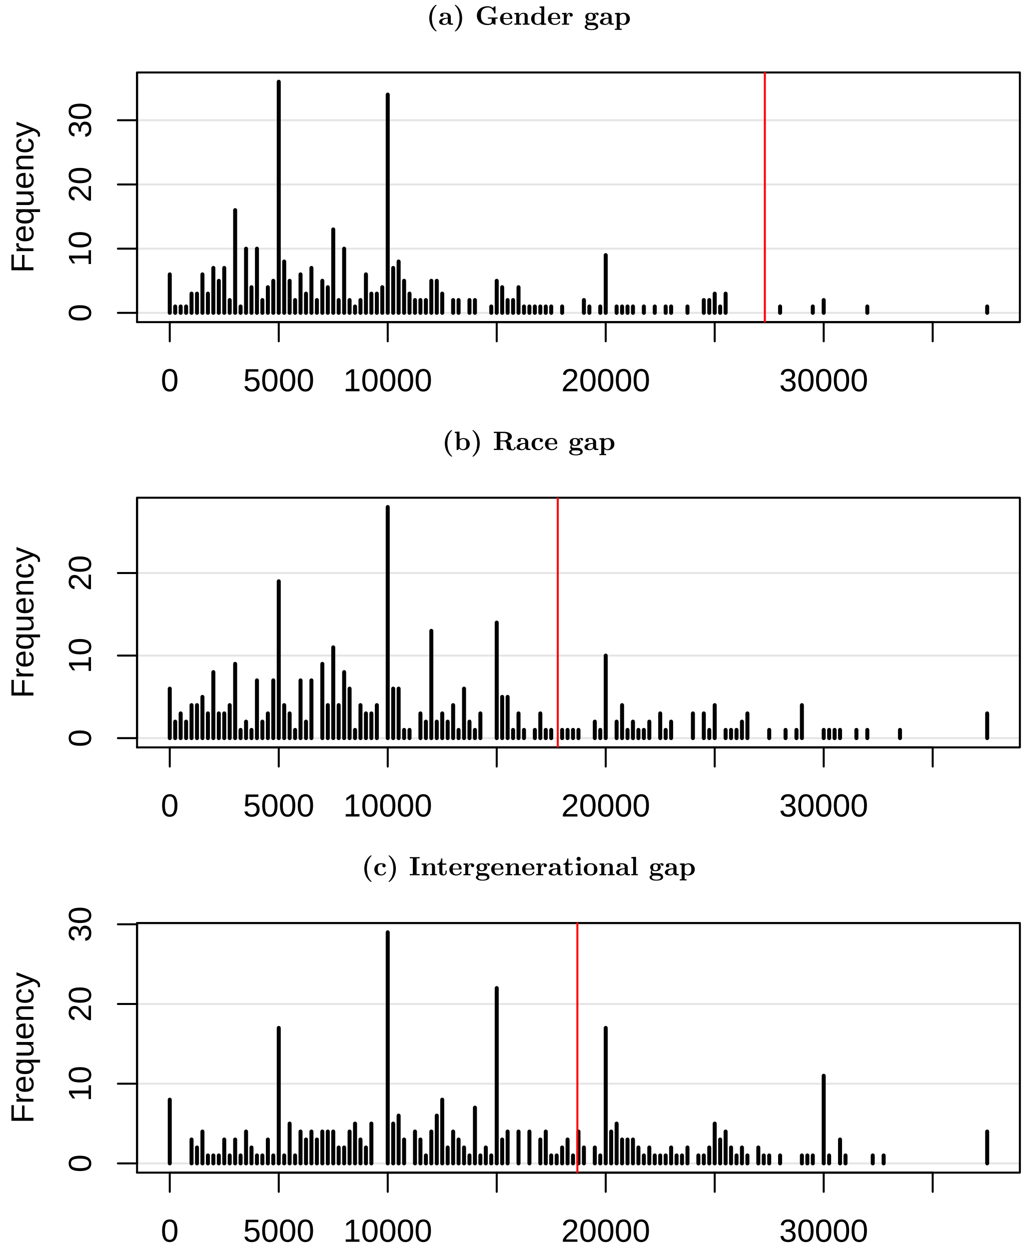

Note: Black bars indicate the distribution of income gap guesses before treatment. Red lines indicate actual income differences based on PSID data; gender gap, US$27,300; race gap, US$17,800; intergenerational gap, US$18,700 (Becker, 2017).

Appendix A1 show that the results hold if redistribution preferences are treated as a continuous variables.

coefficients can be easily interpreted, as percentage point changes in the probability of agreeing with redistribution.

Model results of the immediate treatment effects in the first survey are presented in Table 2. Model 1 is a null model that includes only an intercept and a dummy for the round in which the respondents was recruited (i.e. May or June 2016). Model 2 adds respondents' prior knowledge about income gaps (in US$10,000). Although the respective coefficient points in a positive direction, it is not significant. This is not surprising as the effect of prior knowledge should depend on whether one receives accurate information or not. Model 3 adds a dummy variable that indicates whether a person has been treated or not. Thus, the respective estimate indicates the percentage change in the predicted probability of agreeing with redistribution for the average respondent. The average respondent underestimates the income gaps, and the parameter estimate points as expected in a positive direction. However, the parameter falls short of being statistically significant. This changes in model 4 that adds an interaction term between treatment status and prior knowledge. In this model, the parameter estimate of prior knowledge reflects the relationship between the income gap guesses and agreement with redistribution in the control group. In the control group, which received no information about the accurate gaps, this relationship is as expected positive. Furthermore, the treatment indicator reflects the effect of the informational treatment on a person who maximally underestimated the income gaps, i.e. US$0. This effect is large and statistically significant. Finally, the estimate of the

Table 2: Estimation of Treatment Effects on Agreement with Redistribution (Initial Survey).

	Dependent variable:				
	Redistribution (Agreement)				
	(1)	(2)	(3)	(4)	(5)
Treatment(Info)			0.095	0.347*	0.308*
			(0.051)	(0.107)	(0.106)
Prior(Gaps)		0.071	0.073	0.179*	0.157*
		(0.042)	(0.042)	(0.058)	(0.057)
Treatment(Info)*Prior(Gaps)				−0.220*	−0.199*
				(0.083)	(0.082)
Constant	0.568*	0.489*	0.436*	0.315*	0.664*
	(0.037)	(0.059)	(0.066)	(0.079)	(0.169)
Round dummy	✓	✓	✓	✓	✓
Controls					✓
Observations	371	371	371	371	370
R^2	0.001	0.008	0.018	0.036	0.098
Adjusted R^2	−0.002	0.003	0.010	0.026	0.065

Note: The controls included in model 5 are dummy variables for gender, race, education, employment status, parental education, and continuous variables for age, income, number of children as well as household members.

*p<0.05

Figure 3: Marginal Treatment Effect on Agreement with Redistribution (Initial Survey).

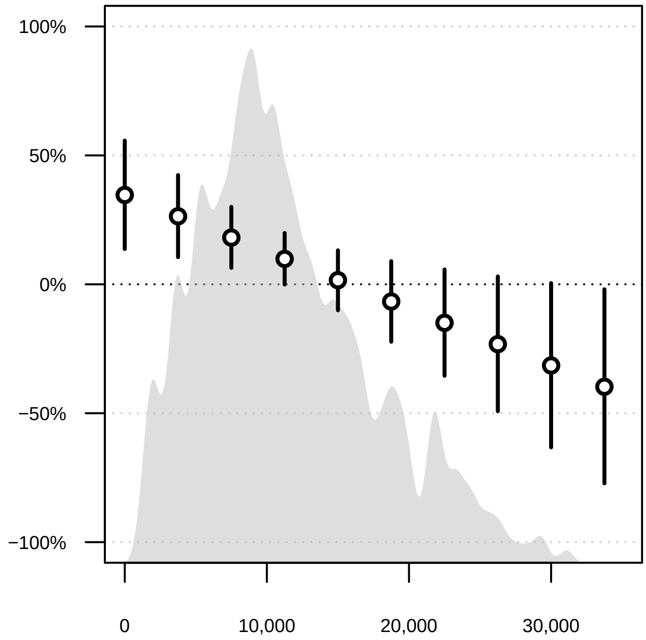

Note: Change in predicted probability of agreeing with redistribution, based on model 4 (Table 2). Confidence intervals (95%) based on bootstrapped model parameters (N=100,000). Grey shading indicates distribution of prior knowledge (kernel density estimation, bandwidth=500).

interaction term shows that this effect decreases the less strongly a person underestimated the gaps. These findings are robust to the inclusion of socio-demographic controls (Model 5).

Based on model 4, I calculate the predicted change in the probability of agreeing with redistribution conditional on a person's prior knowledge. The predicted probabilities are displayed in Figure 3. Those who most severely underestimated the income gaps are most likely to change their stated preference in favor of redistribution. The effect diminishes among those who are closer to the accurate gaps. Those who thought the gaps to be in the proximity of zero become over 30 percentage points more likely to agree with redistribution. Respondents with an average prior guess of US$10,000 still become about 12.5 percentage points more likely to agree with redistribution if confronted with accurate information. There is even some indication that respondents who overestimated the gaps become less likely to agree with redistribution when learning about their accurate size. However, as there are very few respondent who overestimated the gaps this might also be the result of extrapolation from the linear relationship among those who underestimate gaps.

Above analysis posited a linear relationship between prior knowledge and the strength of the treatment effect. Other specifications would require stronger assumptions about how individuals process information. To determine whether positing a linear relationship is warranted I estimate separate regression models for each decile of the *Prior(Gaps)* distribution. Figure 4 displays the treatment effects estimated for each decile. Confidence intervals are wide as the number of observations for each regression is only one tenth of the total sample. The figure attests to clear deviations from a perfectly linear, or even monotonous, relationship. In particular, the treatment effect falls off among those in the lowest decile, whose average prior guess is below US$4,000. One explanation might be that respondents with such low guesses might find information on the actual extent of income gaps hard to believe. Another deviation from a linear relationship is the steep decline in the treatment effect between the fourth and fifth decile. The difference between the two deciles accounts for much of the decline observed over all deciles.

Another important aspect of Figure 4 is that the treatment effect for the tenth decile strongly points into a negative direction. These respondents become more likely to disagree with redistribution when confronted with factual information. This is in fact what would be expected. With average prior guesses of US$20,000 or higher, these respondents overestimated the actual size of the gaps. Therefore, confronting them with factual information should reduce their concern about income gaps and hence demand for redistribution. This corroborates the findings presented in Figure 3 and affirms that the changing sign of the treatment effect is not driven by extrapolation. Overall, the separate regressions attest to a decline in the treatment effect that is sufficiently steady to assume

Figure 4: Treatment Effect on Agreement with Redistribution for Prior Knowledge Deciles (Initial Survey).

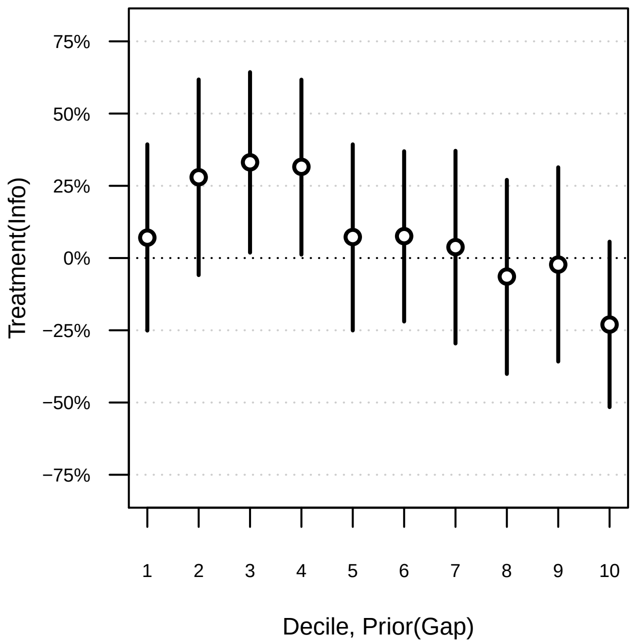

Note: Vertical axis indicates estimated parameter of treatment status, i.e. predicted change in probability of agreeing with redistribution (percentage points). Horizontal axis indicates Prior(Gap) decile based on which OLS regression models were estimated. Models include treatment status and round dummy. Confidence intervals (95%) based on bootstrapped model parameters (N=100,000).

a linear relationship between prior knowledge and the effect of factual information.

Two further robustness checks of the findings presented so far are included in Appendix A1. First, the results are the same if the dependent variable is treated as a continuous rather than as a dummy variable (see Table A2). And second, the results are not driven by any income gap in particular. To affirm this, I estimate models 4 and 5 presented in Table 2 for each income gap separately. This robustness check shows that for each income gap, the more respondents underestimated the respective gap the more likely they become to agree with redistribution (see Table A1). Only for the interaction between the treatment on prior knowledge of the race gap, the parameter coefficient is not statistically significant. However, it still points in the expected direction. Overall, the treatment effect is robust to alternative specifications of the dependent variable as well as the independent variables.

Follow-up Survey

After over one year respondents of the initial survey were contacted to participate in a follow-up. About 40% of the original respondents completed the second survey. The attrition analysis in Appendix A2 indicates no systematic differences between the experimental conditions and thus no threats to the initial randomization. Of course, such analyses are limited to observable characteristics.

Results. To determine whether the treatment in the prior year had any lasting effect on support for redistribution, I estimate the same set of models as for the first survey's analysis. The model results of the second survey (Table 3) show a striking resemblance with those of the first (Table 2). The long-term effect of information about income gaps is still strongest among those who underestimated them most strongly and is less strong among those whose guesses where more accurate. In fact, the coefficients in most models here indicate somewhat stronger effects than documented for the initial survey. At the same time most standard errors are larger now. However, this is easily explained by the smaller sample size, and more importantly, the relevant coefficients are still statistically significant. Again, this finding is robust to the inclusion of the set of control variables (model 6).

Although the analysis above revealed no evidence of any attrition bias based on observables, it is possible to estimate regression models that account for the induced imbalances. Therefore, one first has to determine the probability of all original respondents to participate in the second survey. This is done by estimating a model that predicts the probability of participating in the second survey based on covariates. In a second step, one applies the inverse of the predicted probabilities as weights in the analysis of treatment effects. This process gives more weight to those respondents in the second

Table 3: Estimation of Treatment Effects on Agreement with Redistribution (Follow-up Survey).

	Dependent variable:					
	Redistribution (Agreement)					
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment(Info)			0.012	0.405*	0.586*	0.401*
			(0.078)	(0.163)	(0.159)	(0.165)
Prior(Gaps)		0.171*	0.171*	0.323*	0.362*	0.273*
		(0.064)	(0.065)	(0.084)	(0.078)	(0.087)
Treatment(Info)*Prior(Gaps)				−0.346*	−0.453*	−0.334*
				(0.128)	(0.120)	(0.129)
Constant	0.671*	0.482*	0.475*	0.291*	0.218	0.416
	(0.056)	(0.090)	(0.101)	(0.120)	(0.116)	(0.275)
Round dummy	✓	✓	✓	✓	✓	✓
Reweighted					✓	
Controls						✓
Observations	144	144	144	144	144	144
R^2	0.00002	0.048	0.048	0.096	0.146	0.174
Adjusted R^2	−0.007	0.034	0.027	0.070	0.121	0.092

Note: Inverse probability weighting in model 5 and controls in model 6 based on the same variables, i.e. dummy variables for gender, race, education, employment status, parental education, and continuous variables for age, income, number of children as well as household members. $^*p<0.05$

survey that have similarities with those who dropped out before, thus accounting even for minor covariate imbalances Gerber and Green (2012). Model 5 in Table 2 displays the results of a regression analysis that applies inverse probability weights, which have been determined based on the same set of controls as in model 6.[11] The results shows that correcting for covariate imbalances strengthens the evidence in favor of a treatment effect and its interaction with prior knowledge.

For reasons discussed above, respondents' answer to the redistribution question is coded as a binary variable in the analyses presented here. However, as a robustness check I estimate the same models with redistributive preferences as continuous dependent variable, ranging from 1 (Strongly disagree) to 7 (Strongly agree). Although, the interaction effect loses statistical significance for one of the model specifications, the robustness check reveals the same patterns and thus corroborates the results presented here (see Appendix A1, Table A2).

Manipulation check. An advantage of re-surveying respondents is the possibility to determine whether the treatment had a lasting effect on their knowledge about income gaps. Therefore, the second survey again asked respondents about their guess of the three

[11]Table A3 in Appendix A1 shows how the results of the model based on which inverse probability weights have been computed.

Table 4: Estimation of Treatment Effects on Knowledge about Income Gaps (Follow-up Survey).

	Dependent variable:					
	Income Gaps (Average Guess)					
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment(Info)			0.040	0.398*	0.463*	0.416*
			(0.092)	(0.196)	(0.195)	(0.199)
Prior(Gaps)		0.496*	0.497*	0.635*	0.600*	0.635*
		(0.076)	(0.077)	(0.101)	(0.095)	(0.105)
Treatment(Info)*Prior(Gaps)				−0.316*	−0.320*	−0.340*
				(0.153)	(0.147)	(0.156)
Constant	1.090*	0.540*	0.517*	0.349*	0.342*	0.397
	(0.074)	(0.107)	(0.120)	(0.144)	(0.143)	(0.333)
Round dummy	✓	✓	✓	✓	✓	✓
Reweighted					✓	
Controls						✓
Observations	144	144	144	144	144	144
R^2	0.020	0.246	0.247	0.269	0.262	0.321
Adjusted R^2	0.013	0.235	0.231	0.248	0.240	0.254

Note: Inverse probability weighting in model 5 and controls in model 6 based on the same variables, i.e. dummy variables for gender, race, education, employment status, parental education, and continuous variables for age, income, number of children as well as household members. $^{*}p<0.05$

income gaps. Put differently, I check whether the treatment successfully and lastingly manipulated respondents' knowledge. The regression results in Table 4 show that the average respondent continues to underestimate the gaps. This is most clearly visible from model 1 which features only an intercept, which indicates the average guess of all respondents in the follow-up survey. Model 2 adds prior knowledge, that is respondents' guess in the first survey. It shows that prior knowledge and the guess in the follow-up survey are strongly related. This also explains the good model fit for this and the following models. Most importantly, the fully specified model 4 attests to lasting and statistically significant treatment effects in the expected direction. Those who initially underestimated the gaps lastingly correct their guesses upwards. This correction diminishes among respondents with more accurate prior knowledge. The finding is also robust against the inclusion of controls (model 6) and becomes more pronounced if inverse probability weights are applied (model 5).

In sum, the models here show that the information provided to the treated did lead them to correct their knowledge about income gaps. What is more, the revealed coefficient patterns coincide with those of the above model assessing the treatment effects on agreement with redistribution. This strongly supports the argument that the treatment effect on redistribution preferences is the result of updated beliefs/knowledge about income gaps.

Further analyses speak to the robustness of the presented results. As for the initial survey, I analyzed whether the results hold when each income gap is analyzed separately. Table A5 in Appendix A1 shows that, by and large, treatment effects are present for each separate income gap. This is true for both agreement with redistribution and beliefs about income gaps. Only with regards to the intergenerational gap, the effect of information on the accuracy of the guess in the follow-up survey fully disappears; however, the effect on redistributive preferences persists. In another robustness check, the sample is limited to paid participants of the follow-up (N=115). I re-estimate models 4, 5, and 6 of Tables 3 and 4. The point estimates of all parameters are barely different for the limited sample (see Table A6). However, due to wider confidence intervals, which are partly induced by the reduced sample size, some parameter estimates are not statistically significant anymore.

Party-contingent effects

Earlier research suggests that ideological predispositions can affect responses to factual information. Such predispositions correspond strongly to people's partisanship. Therefore, I explore here whether the effect of information about income gaps differs between Republicans and Democrats.

In the second survey, respondents were asked how they position themselves politically, that is as Democrats, Republicans, Independents, or Other. Those falling into one of the two latter categories were further probed whether they lean toward the Republican or Democratic party. For the following analysis, independents and followers of "Other" parties that indicated no leaning were excluded (N=11). Of the remaining respondents those who identified as Democrats or as leaning towards the party are categorized as *Democrat*. Hence, Republicans constitute the reference group in the following analysis. The remaining sample includes 133 respondents, with 90 Democrats and 43 Republicans.

The results above found that even after one year, respondents who strongly underestimated the income gaps remain more likely to support redistribution. Is this effect equally present among Democrats and Republicans? I begin by estimating a model that adds partisanship as a control to the models presented above. Model 1 (Table 5) shows what is widely known, support of governmental redistribution is much more common among Democrats. Democrats who think income gaps are non-existent are 87.1 percentage points more likely to support redistribution. While the interaction with prior knowledge is statistically insignificant, it indicates that the difference between Democrats and Republicans decreases among those who thought the income gaps were larger.

Another important point to note about this model, and the ones that follow, is that it is "harder" to attain statistically significant results. This is due to the loss in degrees of freedom, which follows from the decreased sample size and the additionally estimated coefficients (i.e. partisanship and accompanying interactions). This likely explains why,

Table 5: Estimation of Treatment Effects, by Party ID (Follow-up Survey).

	Dependent variable:							
	Redistribution (Agreement)				Income Gaps (Average Guess)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment(Info)	0.232	0.324	0.392	0.393	0.460*	1.147*	1.241*	1.247*
	(0.122)	(0.220)	(0.221)	(0.227)	(0.198)	(0.356)	(0.364)	(0.371)
Prior(Gaps)	0.318*	0.483*	0.505*	0.535*	0.550*	0.852*	0.907*	0.884*
	(0.109)	(0.148)	(0.143)	(0.157)	(0.177)	(0.239)	(0.235)	(0.256)
Democrat	0.871*	0.934*	0.919*	0.992*	−0.167	0.333	0.409	0.423
	(0.133)	(0.187)	(0.183)	(0.197)	(0.215)	(0.303)	(0.301)	(0.322)
Tr.(Info)*Pr.(Gaps)	−0.225*	−0.461*	−0.540*	−0.519*	−0.355*	−0.895*	−0.981*	−0.960*
	(0.096)	(0.191)	(0.189)	(0.198)	(0.155)	(0.309)	(0.312)	(0.324)
Tr.(Info)*Dem.		−0.080	−0.085	−0.124		−0.985*	−1.040*	−1.116*
		(0.264)	(0.263)	(0.276)		(0.427)	(0.434)	(0.452)
Pr.(Gaps)*Dem.	−0.190	−0.368*	−0.372*	−0.440*	0.145	−0.247	−0.336	−0.304
	(0.110)	(0.164)	(0.158)	(0.174)	(0.178)	(0.265)	(0.261)	(0.285)
Tr.(Info)*Pr.(Gaps)*Dem.		0.286	0.317	0.329		0.741*	0.815*	0.793*
		(0.220)	(0.216)	(0.230)		(0.356)	(0.357)	(0.377)
Constant	−0.074	−0.154	−0.172	−0.336	0.426*	0.062	0.010	0.209
	(0.129)	(0.161)	(0.156)	(0.251)	(0.208)	(0.261)	(0.258)	(0.411)
Round dummy	✓	✓	✓	✓	✓	✓	✓	✓
Reweighted			✓				✓	
Controls				✓				✓
Observations	133	133	135	133	133	133	133	133
R^2	0.534	0.553	0.565	0.594	0.309	0.338	0.328	0.385
Adjusted R^2	0.512	0.524	0.537	0.534	0.276	0.295	0.285	0.294

Note: Inverse probability weighting in model 3 and 7, and controls in model 4 and 8 based on the same variables, i.e. dummy variables for gender, race, education, employment status, parental education, and continuous variables for age, income, number of children as well as household members.

*$p<0.05$

for example in model 1, the coefficient of the treatment indicator–although still pointing in the same, expected direction–is not statistically significant any more. That being said, the interaction with prior knowledge remains statistically significant.

To determine whether Democrats and Republicans responded differently to the treatment, model 2 interacts partisanship with the terms including the treatment indicator. Neither for the main effect nor the interaction with prior knowledge does this reveal any significant results. An F-test of both models further affirms that model 2 cannot improve on model 1 (F=2.567, p=.081). These results do not change when selection probabilities are accounted for (model 3) or if control variables are added (model 4). As such, there are no statistically significant differences in how people of different partisanship respond to information about income gaps. Once again, results do not change if the dependent variable is taken to be continuous (see Table A7).

The picture changes with regards to respondents' knowledge about income gaps. Model 5 shows that knowledge about income gaps is equally persistent among Democrats and

Republicans. Between the two surveys knowledge for neither partisan group shifted in any particular direction (see *Constant* and *Democrat*) nor is there a difference in the consistency of the two guesses (see *Pr.(Gaps)**Dem.*). However, adding interactions with the treatment terms, model 6 shows that the treatment did not last equally strong. The main treatment indicator now represents the treatment effect on a Republican who in the initial survey guessed all income gaps to be zero, i.e. the treatment increased the guess in the follow-up to US$11,470. At the same time this effect was much smaller among Democrats. The respective interaction indicates that a Democrat, other things being equal, increased their guess from 0 to 1620 only. Similarly, the triple interaction with prior knowledge indicates that information was less effective in changing Democrats' knowledge of income gaps. While these findings are largely robust to adjustment of selection probabilities (model 7) and the inclusion of control variables (model 8), F-tests indicate that adding interactions with partisanship do not significantly improve the model fit.[12]

Discussion

The findings presented in this paper come with the same caveats that apply to similar experiments. First of all, even though the respondents in this experiment cover a wide range of socio-demographics they constitute a convenience sample and findings cannot be generalized to the American populace. Therefore, repeating the experiment on different samples or in representative surveys is of utmost importance. Furthermore, Barabas and Jerit (2010) have shown that informational effects are contingent on levels of exposure. Hence, survey experiments usually find stronger effects than more realistic natural or field experiments. Ideally, future experiments on the mechanism revealed in this paper will make use of such designs. Despite these shortcomings, the present paper also overcame one major criticism of earlier survey experiments, the durability of effects. While this finding certainly calls for replications, the fact that treatment effects persisted for well over one year should be encouraging for other scholars interested in the effect of information on individual beliefs and preferences.

Substantively, the mechanism revealed in this paper connects two strains of political economy scholarship. One of them argues that it is beliefs about equality of opportunity or economic fairness that are decisive for redistributive preferences. With the exception of a few works on intergenerational mobility, this scholarship does not link preferences and objective characteristics of the income distribution. The second line of scholarship focuses on how such objective characteristics influence preferences. However, that scholarship has found it difficult to determine what it is about inequality that people reject, unless it is aligned with their material self-interest. This paper argues that income gaps–an

[12] F-test of model 5 and 6: F=2.669 (p=.073).

objective characteristic of the income distribution–serve people as an indication for the presence of unequal opportunities. As such, the paper opens a new avenue to explore how developments in the income distribution and demand for government redistribution relate.

Great care was taken in the experimental design to ascertain that any revealed effect can indeed be attributed to the informational content of the treatment. Still, some questions about the underlying mechanism remain. While I argued that it is a desire for distributive equality of opportunity that underpins the effect, one can also argue that respondents use the information to update other beliefs that are relevant to redistributive preferences. For example, people might form preferences according to the Rawlsian difference principle and use the provided information to update their beliefs about the well-being of underprivileged groups. Alternatively, learning that income gaps are different from what one thought might lead individuals to update beliefs about the average economic standing and thus also their own relative economic standing. While this implies a rather complex mechanism, the possibility of some hidden self-interest cannot be fully excluded. It is up to future research to better discriminate between these mechanisms.

Another important insight of this paper is that the effect of information about income gaps on redistributive preferences is constant across the ideological divide. Earlier experimental studies on inequality and social mobility have shown that informational effects are often limited to one side of the political spectrum. Although it is widely assumed that Democrats are more likely to subscribe to a distributive understanding of equality of opportunity and income gaps should thus be more relevant to their preference formation, this was not confirmed by the present study. Instead, the findings presented here suggest that information about income gaps and appeals to distributive equality of opportunity might be one of the few topics in contemporary American politics that are not fraught with polarization.

In addition to replications and extensions of the presented experiment, it is important to study how information about the income distribution spreads in the real world. An interesting starting point is work by Iversen and Soskice (2015) who argue that inequality and lack of information about it are reinforcing. They show that increases in inequality are associated with institutional change, like decreasing union density and access to education, which simultaneously undermine the availability of political information to the poor. It is thinkable that similar dynamics are at work with regards to income gaps. In particular, income gaps can limit the resources disadvantaged groups have at their avail to contest such gaps and inform the public about them.

References

Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso (2017). Intergenerational Mobility and Support for Redistribution. *NBER Working Paper 23027*.

Balafoutas, Loukas, Martin G. Kocher, Louis Putterman, and Matthias Sutter (2013, May). Equality, equity and incentives: An experiment. *European Economic Review 60*, 32–51.

Barabas, Jason and Jennifer Jerit (2010, May). Are Survey Experiments Externally Valid? *American Political Science Review 104*(02), 226–242.

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

Becker, Bastian (2017). *Opportunity Matters: Economic Inequality and Redistribution in Western Democracies*. Budapest, Hungary: Central European University. Unpublished dissertation manuscript.

Berinsky, Adam. J., Gregory. A. Huber, and Gabriel. S. Lenz (2012, July). Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk. *Political Analysis 20*(3), 351–368.

Camerer, Colin F., Anna Dreber, Eskil Forsell, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Johan Almenberg, Adam Altmejd, Taizan Chan, Emma Heikensten, Felix Holzmeister, Taisuke Imai, Siri Isaksson, Gideon Nave, Thomas Pfeiffer, Michael Razen, and Hang Wu (2016, March). Evaluating replicability of laboratory experiments in economics. *Science 351*(6280), 1433–1436.

Clifford, Scott, Ryan M Jewell, and Philip D Waggoner (2015, October). Are samples drawn from Mechanical Turk valid for research on political ideology? *Research & Politics 2*(4).

Collaboration, Open Science (2015, August). Estimating the reproducibility of psychological science. *Science 349*(6251), aac4716.

Cruces, Guillermo, Ricardo Pérez Truglia, and Martin Tetaz (2012). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics*.

Dimick, Matthew, David Rueda, and Daniel Stegmueller (2017). The Altruistic Rich? Inequality and other-regarding preferences for redistribution. *Quarterly Journal of Political Science 11*(4), 385–439.

Dion, Michelle L. and Vicki Birchfield (2010, June). Economic Development, Income Inequality, and Preferences for Redistribution. *International Studies Quarterly 54*(2), 315–334.

Evans, Geoffrey and Robert Andersen (2006, February). The Political Conditioning of Economic Perceptions. *Journal of Politics 68*(1), 194–207.

Evans, M. D. R. and Jonathan Kelley (2004, March). Subjective Social Location: Data From 21 Nations. *International Journal of Public Opinion Research 16*(1), 3–38.

Fehr, Ernst and Klaus M. Schmidt (1999, August). A Theory of Fairness, Competition, and Cooperation. *The Quarterly Journal of Economics 114*(3), 817–868.

Fernández-Albertos, José and Alexander Kuo (2015, December). Income Perception, Information, and Progressive Taxation: Evidence from a Survey Experiment. *Political Science Research and Methods*, 1–28.

Finseraas, Henning (2009, March). Income Inequality and Demand for Redistribution: A Multilevel Analysis of European Public Opinion. *Scandinavian Political Studies 32*(1), 94–119.

Fong, Christina (2001). Social preferences, self-interest, and the demand for redistribution. *Journal of Public Economics 82*(2), 225–246.

Gaines, Brian J., James H. Kuklinski, and Paul J. Quirk (2007). The Logic of the Survey Experiment Reexamined. *Political Analysis 15*(01), 1–20.

Gerber, Alan S. and Donald P. Green (2012, May). *Field Experiments: Design, Analysis, and Interpretation.* New York: W. W. Norton & Company.

Griffiths, Thomas L., Charles Kemp, and Joshua B. Tenenbaum (2008). Bayesian Models of Cognition. In R. Sun (Ed.), *The Cambridge Handbook of Computational Psychology*, pp. 59–100. Cambridge University Press.

Haggard, Stephan, Robert R. Kaufman, and James D. Long (2013, April). Income, Occupation, and Preferences for Redistribution in the Developing World. *Studies in Comparative International Development 48*(2), 113–140.

Hainmueller, Jens and Daniel J. Hopkins (2015, July). The Hidden American Immigration Consensus: A Conjoint Analysis of Attitudes toward Immigrants. *American Journal of Political Science 59*(3), 529–548.

Hopkins, Daniel J., John Sides, and Jack Citrin (2016, June). The Muted Consequences of Correct Information About Immigration. SSRN Scholarly Paper ID 2798622, Social Science Research Network, Rochester, NY.

Iversen, Torben and David Soskice (2015). Information, Inequality and Mass Polarization: Ideology in Advanced Democracies. *Comparative Political Studies*.

Jaime-Castillo, Antonio M. and Ildefonso Marqués-Perales (2014, July). Beliefs about Social Fluidity and Preferences for Social Policies. *Journal of Social Policy 43*(03), 615–633.

Karadja, Mounir, Johanna Mollerstrom, and David Seim (2016, June). Richer (and Holier) Than Thou? The Effect of Relative Income Improvements on Demand for Redistribution. *The Review of Economics and Statistics 99*(2), 201–212.

Krawczyk, Michał (2010, February). A glimpse through the veil of ignorance: Equality of opportunity and support for redistribution. *Journal of Public Economics 94*(1-2), 131–141.

Kuklinski, James H., Paul J. Quirk, Jennifer Jerit, David Schwieder, and Robert F. Rich (2000). Misinformation and the currency of democratic citizenship. *Journal of Politics 62*(3), 790–816.

Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva (2015, April). How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments. *American Economic Review 105*(4), 1478–1508.

Lübker, Malte (2007, January). Inequality and the demand for redistribution: are the assumptions of the new growth theory valid? *Socio-Economic Review 5*(1), 117–148.

Lenz, Gabriel S. (2009). Learning and Opinion Change, Not Priming: Reconsidering the Priming Hypothesis. *American Journal of Political Science 53*(4), 821–837.

Lindert, Peter H. and Jeffrey G. Williamson (2016). *Unequal gains: American growth and inequality since 1700.* The Princeton economic history of the Western world. Princeton: Princeton University Press.

Linos, Katerina and Martin West (2003, September). Self-interest, Social Beliefs, and Attitudes to Redistribution. Re-addressing the Issue of Cross-national Variation. *European Sociological Review 19*(4), 393–409.

Loveless, Matthew and Stephen Whitefield (2011, March). Being unequal and seeing inequality: Explaining the political significance of social inequality in new market democracies. *European Journal of Political Research 50*(2), 239–266.

McCall, Leslie (2013). *The Undeserving Rich: American Beliefs about Inequality, Opportunity, and Redistribution.* Cambridge University Press.

McCall, Leslie and Lane Kenworthy (2009, September). Americans' Social Policy Preferences in the Era of Rising Inequality. *Perspectives on Politics 7*(03), 459–484.

McDermott, Rose (2002). Experimental methodology in political science. *Political Analysis 10*(4), 325–342.

Meltzer, Allan H. and Scott F. Richard (1981). A rational theory of the size of government. *The Journal of Political Economy 89*(5), 914–927.

Milanovic, Branko (2000, September). The median-voter hypothesis, income inequality, and income redistribution: an empirical test with the required data. *European Journal of Political Economy 16*(3), 367–410.

Mullinix, Kevin J., Thomas J. Leeper, James N. Druckman, and Jeremy Freese (2015). The Generalizability of Survey Experiments. *Journal of Experimental Political Science 2*(02), 109–138.

Norton, Michael I. and Dan Ariely (2011, January). Building a Better America - One Wealth Quintile at a Time. *Perspectives on Psychological Science 6*(1), 9–12.

Oaxaca, Ronald (1973). Male-Female Wage Differentials in Urban Labor Markets. *International Economic Review 14*(3), 693–709.

Osberg, Lars and Timothy Smeeding (2006, June). "Fair" Inequality? Attitudes toward Pay Differentials: The United States in Comparative Perspective. *American Sociological Review 71*(3), 450–473.

Roemer, John E (1998). *Equality of opportunity.* Cambridge, Mass.: Harvard University Press.

Romer, Thomas (1975, February). Individual welfare, majority voting, and the properties of a linear income tax. *Journal of Public Economics 4*(2), 163–185.

Rueda, David, Daniel Stegmueller, and Timo Idema (2014). The effects of income expectations on redistribution preferences in Western Europe. Washington, DC. American Political Science Association.

Sauermann, Jan and André Kaiser (2010). Taking Others into Account: Self-Interest and Fairness in Majority Decision Making. *American Journal of Political Science 54*(3), 667–685.

Schmidt-Catran, Alexander W. (2016, January). Economic inequality and public demand for redistribution: combining cross-sectional and longitudinal evidence. *Socio-Economic Review 14*(1), 119–140.

Sides, John (2016, May). Stories or Science? Facts, Frames, and Policy Attitudes. *American Politics Research 44*(3), 387–414.

Stewart, Neil, Christoph Ungemach, Adam JL Harris, Daniel M. Bartels, Ben R. Newell, Gabriele Paolacci, and Jesse Chandler (2015). The average laboratory samples a population of 7,300 Amazon Mechanical Turk workers. *Judgment and Decision Making 10*(5), 479.

Trump, Kris-Stella (2017, January). Income Inequality Influences Perceptions of Legitimate Income Differences. *British Journal of Political Science*, 1–24.

Online Appendix

A1 Robustness Checks

Table A1: Robustness of Treatment Effects for Each Income Gap (Initial Survey).

	Dependent variable:					
	Redistribution (Agreement)					
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment(Info)	0.280*	0.237*	0.221*	0.197*	0.319*	0.282*
	(0.087)	(0.087)	(0.088)	(0.087)	(0.101)	(0.100)
Prior(Gender Gap)	0.173*	0.138*				
	(0.055)	(0.055)				
Prior(Race Gap)			0.149*	0.136*		
			(0.048)	(0.047)		
Prior(Intergenerational Gap)					0.065	0.058
					(0.043)	(0.043)
Tr.(Info)*Prior(Gender Gap)	−0.201*	−0.169*				
	(0.079)	(0.078)				
Tr.(Info)*Prior(Race Gap)			−0.118	−0.110		
			(0.064)	(0.063)		
Tr.(Info)*Prior(Intergen. Gap)					−0.156*	−0.141*
					(0.062)	(0.062)
Constant	0.359*	0.759*	0.359*	0.750*	0.431*	0.748*
	(0.068)	(0.164)	(0.068)	(0.159)	(0.075)	(0.171)
Round dummy	✓	✓	✓	✓	✓	✓
Controls		✓		✓		✓
Observations	369	368	370	369	367	366
R^2	0.036	0.096	0.036	0.100	0.028	0.093
Adjusted R^2	0.026	0.063	0.026	0.067	0.017	0.059

Note: Separate re-estimations of models 4 and 5 (Table 2) for each income gap. Controls in odd-numbered models are dummy variables for gender, race, education, employment status, parental education, and continuous variables for age, income, number of children as well as household members.

*p<0.05

Table A2: Robustness of Treatment Effects on Agreement with Redistribution (Initial Survey).

	Dependent variable:				
	Redistribution (1-7)				
	(1)	(2)	(3)	(4)	(5)
Treatment(Info)			0.228	1.162*	0.896*
			(0.222)	(0.467)	(0.455)
Prior(Gap)		0.338	0.343	0.738*	0.619*
		(0.181)	(0.181)	(0.251)	(0.245)
Treatment(Info)*Prior(Gap)				−0.819*	−0.644
				(0.361)	(0.351)
Constant	4.437*	4.060*	3.934*	3.483*	6.021*
	(0.158)	(0.257)	(0.285)	(0.346)	(0.722)
Round dummy	✓	✓	✓	✓	✓
Controls					✓
Observations	371	371	371	371	370
R^2	0.003	0.012	0.015	0.029	0.122
Adjusted R^2	0.0004	0.007	0.007	0.018	0.090

Note: Continuous dependent variable ranging from 1 (Strongly disagree) to 7 (Strongly agree), otherwise models equivalent to Table 2. Controls in model 5 are dummy variables for gender, race, education, employment status, parental education, and continuous variables for age, income, number of children as well as household members.

*p<0.05

Table A3: Predicting probabilities for second survey participation

	Dependent variable:
	Participation, Follow-up Survey
Treatment(Info)	−0.241
	(1.776)
Redistribution (Survey 1)	−0.034
	(0.083)
Prior(Gap)	−0.086
	(0.275)
Round	0.198
	(0.335)
Duration (Survey 1)	−0.056
	(0.038)
Male	−0.670
	(0.371)
White	−0.181
	(0.469)
University-educated parent	−0.013
	(0.387)
Full-time employment	0.730
	(0.426)
Children	0.200
	(0.202)
Household size	−0.155
	(0.176)
Income	−0.091
	(0.053)
University degree	0.423
	(0.397)
Age	0.043*
	(0.020)
Treatment(Info)*Redistribution(Survey 1)	−0.069
	(0.114)
Treatment(Info)*Prior(Gap)	−0.108
	(0.388)
Treatment(Info)*Round	−0.190
	(0.463)
Treatment(Info)*Duration (Survey 1)	0.057
	(0.049)
Treatment(Info)*Male	0.662
	(0.498)
Treatment(Info)*White	−0.220
	(0.634)
Treatment(Info)*Univ.-educated parent	0.047
	(0.514)
Treatment(Info)*Full-time employment	−0.885
	(0.562)
Treatment(Info)*Children	−0.271
	(0.292)
Treatment(Info)*Household size	0.109
	(0.224)
Treatment(Info)*Income	0.068
	(0.081)
Treatment(Info)*University degree	−0.089
	(0.528)
Treatment(Info)*Age	0.002
	(0.026)
Constant	−0.762
	(1.402)
Observations	370
Log Likelihood	−230.061
Akaike Inf. Crit.	516.122
Note: Binary dependent variable indicating participation in follow-up survey. Parameter estimated using generalized linear model with logit link function.	*p<0.05

Table A4: Robustness of Treatment Effects on Agreement with Redistribution (Follow-up Survey).

	Dependent variable:					
	Redistribution (1-7)					
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment(Info)			0.029	1.873*	0.586*	1.731*
			(0.343)	(0.719)	(0.159)	(0.732)
Prior(Gap)		0.663*	0.664*	1.378*	0.362*	1.222*
		(0.284)	(0.286)	(0.372)	(0.078)	(0.386)
Treatment(Info)*Prior(Gap)				−1.626*	−0.453*	−1.512*
				(0.562)	(0.120)	(0.573)
Constant	4.614*	3.878*	3.861*	2.996*	0.218	3.172*
	(0.248)	(0.399)	(0.447)	(0.528)	(0.116)	(1.224)
Round dummy	✓	✓	✓	✓	✓	✓
Re-weighted					✓	
Controls						✓
Observations	144	144	144	144	144	144
R^2	0.003	0.040	0.040	0.095	0.146	0.158
Adjusted R^2	−0.004	0.027	0.020	0.069	0.121	0.074

Note: Continuous dependent variable ranging from 1 (Strongly disagree) to 7 (Strongly agree), otherwise models equivalent to Table 3. Inverse probability weighting in model 5 and controls in model 6 based on the same variables, i.e. dummy variables for gender, race, education, employment status, parental education, and continuous variables for age, income, number of children as well as household members.

$^{*}p<0.05$

Table A5: Robustness of Treatment Effects for Each Income Gap (Follow-up Survey).

	Dependent variable:					
	Redistribution (Agreement)			Income Gaps Guess (Follow-up Survey)		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment(Info)	0.217	0.212	0.322*	0.378*	0.400*	0.025
	(0.129)	(0.131)	(0.158)	(0.177)	(0.196)	(0.259)
Prior(Gender Gap)	0.267*			0.548*		
	(0.077)			(0.105)		
Prior(Race Gap)		0.239*			0.709*	
		(0.070)			(0.103)	
Prior(Intergenerational Gap)			0.159*			0.402*
			(0.068)			(0.109)
Tr(Info)*Prior(Gender Gap)	−0.219			−0.344*		
	(0.115)			(0.157)		
Tr(Info)*Prior(Race Gap)		−0.207*			−0.356*	
		(0.096)			(0.142)	
Tr(Info)*Prior(Intergen. Gap)			−0.229*			0.045
			(0.101)			(0.166)
Constant	0.420*	0.426*	0.463*	0.307*	0.357*	0.692*
	(0.103)	(0.103)	(0.116)	(0.142)	(0.154)	(0.188)
Round dummy	✓	✓	✓	✓	✓	✓
Observations	145	145	144	142	139	140
R^2	0.081	0.078	0.044	0.204	0.322	0.171
Adjusted R^2	0.055	0.051	0.016	0.181	0.302	0.147

Note: Separate re-estimations of models 4 in Table 3 and 4 for each income gap. In models 4-6, the dependent variable is the guess of the gender, race, respectively intergenerational income gap in the follow-up survey.

*p<0.05

Table A6: Robustness of Treatment Effects (Follow-up Survey, Sample: Paid Respondents)

	Dependent variable:					
	Redistribution(Agreement)			Income Gaps (Average Guess)		
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment(Info)	0.428*	0.432*	0.595*	0.407	0.425	0.469*
	(0.184)	(0.184)	(0.176)	(0.213)	(0.218)	(0.207)
Prior(Gaps)	0.360*	0.344*	0.382*	0.567*	0.588*	0.512*
	(0.101)	(0.101)	(0.090)	(0.117)	(0.120)	(0.106)
Tr.(Info)*Prior(Gaps)	−0.366*	−0.374*	−0.468*	−0.276	−0.323	−0.282
	(0.146)	(0.147)	(0.137)	(0.170)	(0.175)	(0.161)
Constant	0.237	0.147	0.186	0.422*	0.391	0.457*
	(0.137)	(0.311)	(0.132)	(0.159)	(0.370)	(0.155)
Round dummy	✓	✓	✓	✓	✓	✓
Controls		✓			✓	
Re-weighted			✓			✓
Observations	115	115	115	115	115	115
R^2	0.104	0.234	0.157	0.223	0.302	0.211
Adjusted R^2	0.072	0.135	0.127	0.195	0.212	0.182

Note: Inverse probability weighting in model 3 and 6, and controls in model 2 and 4 based on the same variables, i.e. dummy variables for gender, race, education, employment status, parental education, and continuous variables for age, income, number of children as well as household members. Sample limited to respondents who participated in the follow-up for pay.

*p<0.05

Table A7: Robustness of Treatment Effect, by Party ID (Follow-up Survey).

	Dependent variable:			
	Redistribution (1-7)			
	(1)	(2)	(3)	(4)
Treatment(Info)	1.150*	1.274	1.591	1.505
	(0.505)	(0.923)	(0.913)	(0.942)
Prior(Gap)	1.307*	1.646*	1.745*	2.228*
	(0.451)	(0.620)	(0.590)	(0.650)
Democrat	3.959*	4.042*	3.932*	4.664*
	(0.548)	(0.785)	(0.756)	(0.816)
Tr.(Info)*Pr.(Gap)	−1.154*	−1.626*	−1.995*	−1.992*
	(0.396)	(0.803)	(0.782)	(0.822)
Tr.(Info)*Dem.		−0.060	−0.201	−0.361
		(1.106)	(1.088)	(1.147)
Pr.(Gap)*Dem.	−0.766	−1.123	−1.116	−1.801*
	(0.456)	(0.687)	(0.655)	(0.724)
Tr.(Info)*Pr.(Gap)*Dem.		0.561	0.784	0.967
		(0.924)	(0.896)	(0.956)
Constant	1.354*	1.217	1.188	−0.313
	(0.532)	(0.676)	(0.647)	(1.043)
Round dummy	✓	✓	✓	✓
Re-weighted			✓	
Controls				✓
Observations	133	133	133	133
R^2	0.595	0.600	0.613	0.644
Adjusted R^2	0.576	0.574	0.588	0.592

Note: Continuous dependent variable ranging from 1 (Strongly disagree) to 7 (Strongly agree), otherwise models equivalent to Table 5. Inverse probability weighting in model 3 and controls in model 4 based on the same variables, i.e. dummy variables for gender, race, education, employment status, parental education, and continuous variables for age, income, number of children as well as household members.

*p<0.05

A2 Attrition Analysis

Experiments that stretch longer time periods unavoidably face attrition, which can lead to bias in the estimation of treatment effects. The high response rate to the second survey is not sufficient to exclude the possibility of such a bias. Therefore, it is important to check for indications of attrition bias. This is done similarly to how researchers check for covariate balance in single-shot experiments. Just as in the case of covariate balance, it is only possible to check for attrition bias based on observables. While the absence of such a bias for observables can increase our confidence in the unbiasedness of results, it is no guarantee.

The most basic source of attrition bias is differential response rates across experimental conditions. Furthermore, it would be worrisome if attrition patterns based on covariates differed between control and treatment group. In order to assess these sources of bias, I run separate regressions for both experimental groups. I begin with an intercept-only model

Table A8: Determinants of Attrition, Regression Results (DV: Participation, Follow-up Survey)

	Control group		Treatment group		Difference (C-T)	
	Estimate	Std. Error	Estimate	Std. Error	Estimate	Std. Error
Intercept	0.407*	0.037	0.387*	0.035	-0.02	0.051
Male	-0.164*	0.073	0	0.071	0.165	0.102
Race: White	0.127	0.089	0.019	0.091	-0.109	0.128
Black	-0.007	0.161	0.121	0.146	0.128	0.217
Other	-0.157	0.1	-0.093	0.109	0.063	0.147
University-educated parent	0.049	0.075	0	0.072	-0.049	0.104
Employment status: Unemployed	-0.143	0.153	-0.041	0.116	0.102	0.192
Full-time	0.076	0.079	-0.035	0.072	-0.111	0.106
Keeping house	-0.164	0.178	-0.012	0.177	0.152	0.251
Other	0.095	0.25	-0.139	0.247	-0.235	0.351
Part-time	-0.133	0.108	0.061	0.105	0.194	0.15
Retired	0.442*	0.202	0.249	0.176	-0.193	0.268
Student	-0.076	0.205	-0.024	0.152	0.052	0.255
Household size	-0.028	0.025	-0.014	0.024	0.014	0.034
Children^{+}	0.043	0.029	0.025	0.032	-0.018	0.044
Income	-0.005	0.009	0.002	0.011	0.007	0.015
Education: Less than high school	-0.409	0.494	-0.389	0.49	0.02	0.696
High school	-0.135	0.076	-0.058	0.074	0.077	0.106
University	0.144	0.075	0.066	0.074	-0.078	0.105
Age	0.009*	0.003	0.009*	0.003	0	0.004
Duration (Survey 1)	-0.003	0.007	0.005	0.007	0.008	0.01
Prior(Gaps)	-0.018	0.058	-0.026	0.06	-0.007	0.083
Redistribution	0	0.017	-0.023	0.017	-0.023	0.024

Note: Results of univariate linear regressions (first row, intercept-only), estimated separately for each covariate, for control (columns 2&3, N=177) and treatment group (columns 4&5, N=194), and jointly with interaction effects between treatment condition and covariate (columns 6&7, N=371). $^{+}$Sample size for these models equals N-1 due to missing response on covariate.

$^{*}p<0.05$

and continue with univariate regressions for the main socio-demographic covariates elicited in the initial survey. Results are shown in Table A8. The response rates of both groups are not exactly the same, 40.7% for the control group and 38.7% for the treatment group. However, as the right-most columns show, the difference is not statistically significant.

There is some evidence, especially in the control group, that respondents who are male, retired, and/or older are more likely to drop out. However, in no case is this pattern significantly different between both groups. There is also no evidence that respondents who took more time for the first survey are more likely to drop out. Finally, it is possible to check for attrition based on the two central variables in this study, respondents' prior knowledge about income gaps and preferences for redistribution. As these variables are constitutive of the causal mechanism explored here, attrition bias would be detrimental. Fortunately, there is no indication of any bias with regards to either variable.