



University of
Zurich^{UZH}

URPP Equality of Opportunity

Beliefs about Inequality and the Nature of Support for Redistribution

*Thomas Epper
Ernst Fehr
Aljosha Henkel
Julien Senn*

Equality of Opportunity Research Series #49
January 2024





**University of
Zurich**^{UZH}

URPP Equality of Opportunity

URPP Equality of Opportunity Discussion Paper Series No. 49, January

Beliefs about Inequality and the Nature of Support for Redistribution

Thomas Epper
IESEG School of Management
thomas.epper@cnrs.fr

Ernst Fehr
University of Zurich
ernst.fehr@econ.uzh.ch

Aljosha Henkel
ETH Zurich
henkel@kof.ethz.ch

Julien Senn
University of Zurich
julien.senn@econ.uzh.ch




The University Research Priority Program “Equality of Opportunity” studies economic and social changes that lead to inequality in society, the consequences of such inequalities, and public policies that foster greater equality of opportunity. We combine the expertise of researchers based at the University of Zurich’s Faculty of Arts and Social Sciences, the Faculty of Business, Economics and Informatics, and the Faculty of Law.

Any opinions expressed in this paper are those of the author(s) and not those of the URPP. Research published in this series may include views on policy, but URPP takes no institutional policy positions.

URPP Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character.

URPP Equality of Opportunity, University of Zurich, Schoenberggasse 1, 8001 Zurich, Switzerland
info@equality.uzh.ch, www.urpp-equality.uzh.ch

Beliefs about Inequality and the Nature of Support for Redistribution*

Aljosha Henkel  Ernst Fehr  Julien Senn  Thomas Epper


January 24, 2024

Abstract

Do beliefs about inequality depend on distributive preferences? What is the joint role of preferences and beliefs about inequality for support for redistribution? We study these questions in a staggered experiment with a representative sample of the Swiss population conducted in the context of a vote on a highly redistributive policy proposal. Our sample comprises a majority of inequality averse subjects, a sizeable group of altruistic subjects, and a minority of predominantly selfish subjects. Irrespective of preference types, individuals vastly overestimate the extent of income inequality. An information intervention successfully corrects these large misperceptions for all types, but essentially does *not* affect aggregate support for redistribution. These results hide, however, important heterogeneity because the effects of beliefs about inequality for demand for redistribution are preference-dependent: only affluent inequality averse individuals, but not the selfish and altruistic ones, significantly reduce their support for redistribution. These findings cast a new light on the seemingly puzzling result that, in the aggregate, large changes in beliefs about inequality often do not translate into changes in demand for redistribution.

Key Words: Social Preferences, Beliefs about Inequality, Preferences for Redistribution, Information, Inequality Aversion

JEL Codes: D31, D72, H23, H24

*The “” symbol indicates that the authors’ names are in certified random order. Ernst Fehr acknowledges the support from the Research Priority Program of the University of Zurich on Equality of Opportunity. Epper: IESEG School of Management, University of Lille, CNRS, UMR 9221 - LEM - Lille Economie Management, F-59000 Lille, France (thomas.epper@cnrs.fr). Fehr: Department of Economics, University of Zurich, Blümlisalpstrasse 10, 8006 Zurich, Switzerland (ernst.fehr@econ.uzh.ch). Henkel: KOF Swiss Economic Institute, ETH Zurich, Leonhardstrasse 21, 8092 Zurich, Switzerland (henkel@kof.ethz.ch). Senn: Department of Economics, University of Zurich, Blümlisalpstrasse 10, 8006 Zurich, Switzerland (julien.senn@econ.uzh.ch).

1 Introduction

Over the last decades, many countries have experienced a substantial increase in income and wealth inequality. This increased concentration of income and wealth has been particularly pronounced at the very top. In the US for example, the share of income captured by the top 1% nearly doubled over the last four decades, from 10.4 percent in 1980 to 19.1 percent in 2020 (World Inequality Database, 2023).

These increasing inequalities have put redistribution again on top of the political agenda and have given rise to civil movements such as the Wallstreet movement (*"We are the 99%"*). Switzerland is no exception: In recent years, several strongly redistributive initiatives have been put to vote. The latest example of such a vote took place in September 2021, when Switzerland voted on the so-called "99% initiative"—a policy proposal aimed at increasing taxes on the capital gains of those at the very top of the income distribution.

But what motivates citizens to support such strongly redistributive proposals? One strand of research has highlighted the role of (biased) beliefs about inequality for support for redistribution (Cruces et al., 2013; Kuziemko et al., 2015; Karadja et al., 2017). Another strand of the literature has underscored the importance of (other-regarding) preferences (Fehr et al., 2022; Fisman et al., 2017; Kerschbamer and Müller, 2020). These two strands of the literature have largely evolved separately. In this paper, we investigate the *joint* role of beliefs about inequality and other-regarding preferences for demand for redistribution. We are particularly interested in understanding whether preferences alter the effects of beliefs about income inequality for support for redistribution, i.e. whether the effects of beliefs about inequality for demand for redistribution are preference-dependent.

This research agenda is motivated by a simple observation: an important premise for beliefs about inequality to matter for demand for redistribution is that individuals care about inequality and the payoffs of others. However, recent field and laboratory evidence indicates that there are important heterogeneities in terms of other-regarding preferences, i.e., not all individuals put the same weight on the distribution of payoffs (see, e.g., Fisman et al., 2017; Kerschbamer and Müller, 2020; Fehr et al.,

2022). This raises several interesting new questions. First, do beliefs about inequality depend on preferences? While previous studies have established that individuals tend to have biased beliefs about inequality (see e.g. Cruces et al., 2013), there may be large heterogeneities in the population. For example, it seems plausible that other-regarding individuals overestimate inequality while selfish individuals underestimate it. Second, can we legitimately expect that beliefs about inequality will shape demand for redistribution to a similar extent across individuals with different preferences? For example, should we expect beliefs about inequality to affect the demand for redistribution of a purely selfish individual? Would providing such an individual with information about inequality affect their demand for redistribution? Would their response to information about inequality be smaller than, say, the response of an inequality averse individual who generally cares about the distribution of payoffs?

We study these questions using a pre-registered online experiment with a representative sample of the Swiss population in the context of the vote on the 99% initiative, a highly redistributive proposal that aimed at increasing taxes for the top 1% by instituting a capital gains' tax. Our experiment consists of two waves, conducted six months apart. In the first wave, we elicit participants' other-regarding preferences using an incentivized money allocation task in which participants have to make distributional choices between themselves and another anonymous participant. In the second wave, we measure participants' beliefs about the income share received by the top 1% of income earners in Switzerland.¹ To assess the causal effects of beliefs about inequality for demand for redistribution, we randomly assign half of the participants to a treatment condition that provides them with credible *factual* information about the share of income received by the top 1% of income-earners in Switzerland. We then elicit all respondents' support for the 99% initiative by allowing them to make a real monetary donation to organizations that either actively campaign in favor of or against this referendum.

This staggered design, which purposefully decouples the elicitation of distributive preferences from the information intervention and the measurement of support

¹In the following, we use "beliefs about (income) inequality" and "perceived (income) inequality" as synonyms for "beliefs about the income share received by the top 1% of the income earners."

for redistribution, allows us to study the causal effects of beliefs about inequality for support for the 99% initiative, and to investigate the extent to which these effects are preference-dependent. In addition, it also allows us to explore whether individuals with different preference types hold fundamentally different beliefs about the extent of inequality, and whether they update these beliefs differently when presented with credible and objective information.

We expect important heterogeneities in subjects' response to our information intervention, depending on subject's preference type and their income. Previous work has provided both theoretical and empirical evidence that other-regarding preferences are an important predictor of political support for redistribution, particularly amongst the most affluent individuals (Fehr & al., 2022).² This is due to the fact that at low incomes it is predominantly selfish motives that fuel support for redistribution, thereby leaving little room for social preferences to add much. In contrast, sufficiently strong social preferences can play a large role at higher incomes by mitigating the strong opposition to redistribution that is due to selfishness. We therefore hypothesize that the effect of beliefs about inequality will be particularly pronounced for inequality averse individuals with an above-median income, relative to selfish individuals with an above-median income.³

We characterize preference heterogeneity in our sample by applying a Bayesian nonparametric clustering algorithm, following the approach discussed in Fehr & al. (2022). Consistent with the findings of Fehr & al. (2022), we also document three distinct types with a clear behavioral interpretation: A large group of predominantly *inequality averse* individuals, a smaller group of *altruistic* individuals, and a minority of predominantly *selfish* individuals. Our sample is therefore very diverse in terms of how respondents weigh other people's payoffs: some individuals are highly sensitive to the distributional consequences of their choices, while others are predominantly self-interested.

We document several novel findings. First, we investigate how individuals *perceive* inequality. We find that our subjects vastly overestimate the extent of income

²We refer to subjects with an income above the median as affluent individuals.

³We pre-registered this conjecture as our main hypothesis.

inequality. While the top 1% of income earners actually receive 12% of the total annual income, the average respondent thinks that the top 1% receives about five times more (average belief: 54.2%). These misperceptions are widespread, with 94% of our sample overestimating inequality, and only 4% underestimating it.

Second, we explore whether these perceptions differ by preference type. We find that these misperceptions are largely orthogonal to respondents' preference type. That is, inequality averse individuals do not have more biased beliefs about inequality than the selfish or the altruistic individuals.

Third, we show that providing respondents with accurate information about the extent of income inequality dramatically reduces these misperceptions, and that all social preferences types update their beliefs to a similar extent. While this informational shock successfully corrects subjects' large misperceptions, it essentially does *not* affect average support for redistribution, consistent with a recent meta-analysis on the causal effects of beliefs about inequality for demand for redistribution (Ciani et al., 2021). This average result hides, however, important heterogeneity. Indeed, as we hypothesized, the downwards shocks in beliefs about inequality generates a large and significant *decrease* in donations in favour of the 99% initiative amongst affluent inequality averse respondents, and it essentially does *not* affect the donations of the remaining subjects. We discuss multiple robustness checks that rule out alternative interpretations to these findings. In particular, we argue that the type-specific nature of our hypotheses, and our results, rule out that our findings are driven by demand effects—a concern against some information provision experiments (Haaland et al., 2023). If our information intervention had generated large demand effects, one would have expected that, regardless of income levels, all preference types show a reduction in the support for redistribution. In contrast, however, only the affluent inequality averse individuals displayed a significant behavioral change—as hypothesized.

Our paper connects to several strands of the literature. First, our paper relates to the large literature on the political economy of demand for redistribution. This literature has identified a list of determinants of support for redistribution, such as beliefs about income mobility (Piketty, 1995; Benabou and Ok, 2001; Benabou and Tirole, 2006; Alesina et al., 2018), beliefs about the causes of success (Fong, 2001; Alesina and

Angeletos, 2005; Almås et al., 2019), and beliefs about income and wealth inequality (Cruces et al., 2013; Karadja et al., 2017; Fehr et al., 2019), among others. More recently, several papers have highlighted the role of preferences, such as equality-altruism tradeoffs (Fisman et al., 2017) and inequality aversion and altruism (Fehr et al., 2022), for the demand for redistribution. We contribute to this literature by investigating the joint role of preferences and beliefs about inequality for the demand for redistribution. We are also the first, to our knowledge, to measure support for redistribution using real monetary donations in the context of a real, upcoming political campaign.

Our paper also connects to the literature that uses information interventions to study subjective beliefs about the economy.⁴ For example, recent papers have used information interventions to study topics as diverse as the role of beliefs about inequality or about the distribution of income for demand for redistribution (Cruces et al., 2013; Karadja et al., 2017; Fehr et al., 2019; Kuziemko et al., 2015), the link between labor market concerns and support for immigration (Haaland and Roth, 2020), the effects of beliefs about public debt for the demand for government spending and taxation (Roth et al., 2021), the relationship between perceptions of existing spending levels on preferences for increased government spending on education (Lergetporer et al., 2018), and the role of beliefs about the size of the gender wage gap for demand for policies intended to reduce it (Settele, 2022). We contribute to this literature by providing novel evidence on the misperceptions of income inequality, and by showing that these misperceptions are widely held amongst individuals characterized by very different preference types. Correcting these misperceptions however only affects support for redistribution of the most affluent inequality averse individuals. To our knowledge, we are the first to document that the effects of beliefs about inequality on demand for redistribution might be preference-dependent. These findings cast a new light on the seemingly puzzling result that, in the aggregate, large changes in beliefs about inequality often do not translate into changes in demand for redistribution, as documented in a recent review of the literature (Ciani et al., 2021).

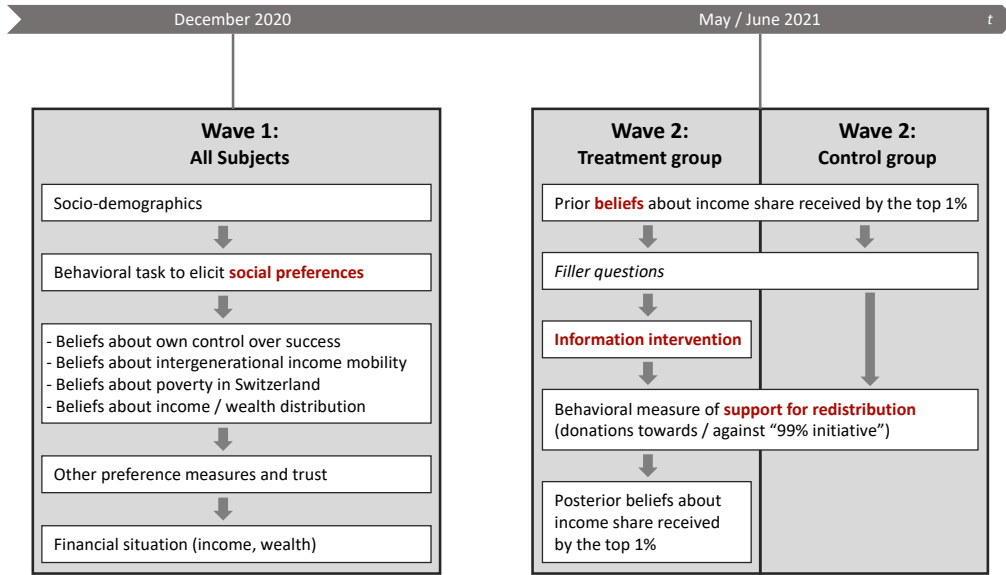
⁴For a recent review of the literature of information interventions, see Haaland et al. (2023).

2 Experimental design

2.1 Overview

Our main aim is to study how beliefs about inequality and distributive preferences jointly predict demand for redistribution. In particular, we are interested in understanding whether individuals with different preference types hold different beliefs about the extent of inequality and whether the causal effect of beliefs on support for redistribution differs across preference types. This endeavour requires the following elements: i) a clean and independent measure of social preferences, ii) a baseline measure of beliefs about inequality, iii) an exogenous shock to beliefs about inequality, and iv) an incentivized measure of demand for redistribution. We approach this task by conducting a staggered experiment with two waves. The key features of our experiment are summarized in Figure 1 below.

Figure 1: Overview of our staggered experiment with two waves



The main goal of the first wave is to measure the distributive preferences of Swiss voters. We also use this wave to collect information on respondents' socio-demographics and on a set of beliefs that have been shown to matter for support for redistribution, such as beliefs about the determinants of individual success, beliefs about intergenerational income mobility, prior (i.e., pre-intervention) beliefs about

poverty, and beliefs about the distributions of income and wealth. In addition, we measure other economic preferences using the preference survey module by Falk et al. (2022). For details on the measurement of these additional variables, see Appendix A.1.

We conduct the second wave six months later. In this second wave, we first measure respondents' prior beliefs about top income inequality in Switzerland (i.e., their beliefs about the share of total income received by the top 1% of income earners). We then exogenously shock these beliefs for half of the respondents by providing them with credible and objective information about the extent of top income inequality in Switzerland. Subsequently, we measure demand for redistribution by allowing all subjects to make a real monetary donation to civic groups that either support or oppose the 99% initiative. Last, we re-elicited treated subjects' beliefs about income inequality in order to assess whether they updated their beliefs.

We provide details on the two waves and on our subject pool in the next subsections.

2.2 Wave 1: Eliciting social preferences

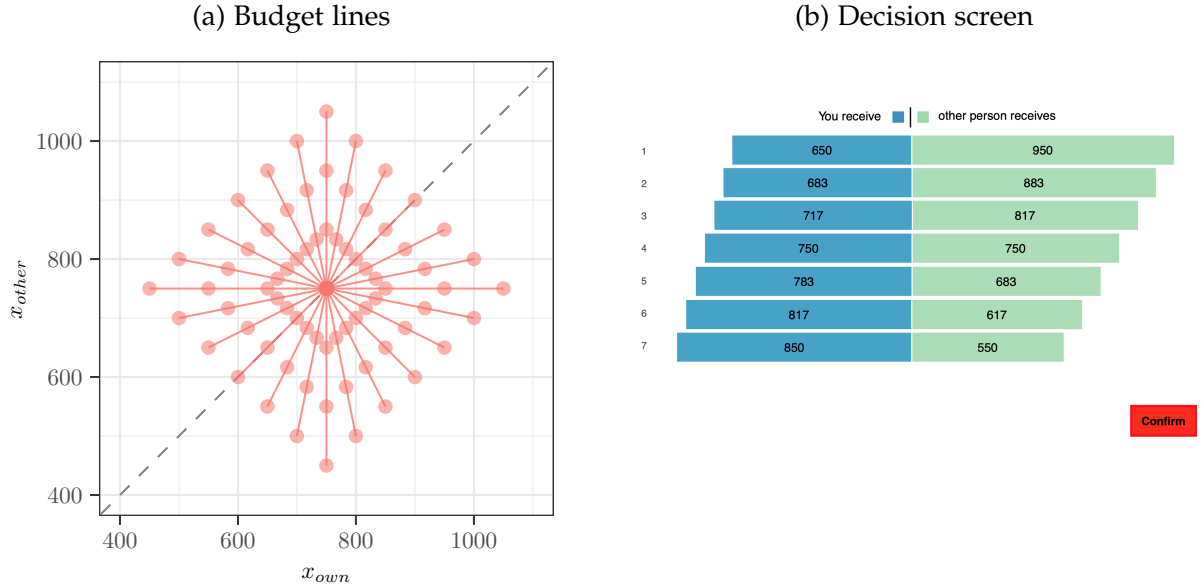
The first wave of our study, which we conducted in December 2020, is aimed at uncovering the distribution of other-regarding preferences among Swiss voters. We approach this task by eliciting respondents' distributional preferences in a set of 12 incentivized money allocation tasks.⁵ In each of these tasks, the participants have to decide how to allocate experimental currency units (ECUs) between themselves and an anonymous other participant of the study.

Figure 2a depicts the various budget lines for which subjects had to make a decision. In some decision situations, the decision maker can give up some of her own payoff to *increase* the payoff of the other. In other decision situations, they can pay to *decrease* the payoff of the other. These different choice situations systematically vary the cost and the joint payoff consequences of redistribution; thereby allowing us to identify a wide range of other-regarding behaviors.

⁵The task used for the clustering and type identification is based on Fehr (r) al. (2022, 2023).

Figure 2b illustrates how a typical choice situation was presented to participants. We represented the available choices numerically and graphically in order to make the trade-offs and the associated payoff implications transparent. There were always seven interpersonal allocations (labeled by 1 to 7) available per choice situation, and all of them were located on a budget line. Each available allocation consisted of a specific distribution of ECUs between the participant (bars labeled by “You receive”) and the other person (bars labeled by “other person receives”). In this example, the slope of the budget line is -2, indicating that for every ECU the decision maker gives up, the other player receives 2 ECUs. Perfect equality in payoffs can be achieved by choosing allocation 4.

Figure 2: Measuring other-regarding preferences with a money allocation task



Note: Figure 2a depicts the various budget lines for which subjects had to make a decision. Figure 2b illustrates how a typical choice situation was presented to participants.

2.3 Wave 2: Beliefs and information provision experiment

2.3.1 Measuring prior beliefs about income inequality

Six months after the first wave, we contacted the *same* subjects to participate in the second part of the study. This wave starts with questions aimed at measuring respondents' prior beliefs about income inequality. More specifically, we elicit subjects' beliefs about the share (in percent) of the total national income that is received by

the top 1% of income-earners in Switzerland. We elicit these beliefs using a two-step approach: We first ask subjects to think about the total income that is received by all the people in Switzerland, and to indicate how much they believe is received by the top 1% of people with the highest incomes by selecting one out of ten possible brackets (e.g., 0%-10%, 11%-20%, 21%-30%, ... , 91%-100%). We then ask respondents to refine their answer by providing a point estimate within the range they chose.⁶

We purposely chose to elicit respondents' beliefs about the *income* share received by the top 1% (as opposed to, for example, the wealth share of the top 1%, or another measure of inequality) because it more closely corresponds to the main purpose of the 99% initiative, which is to increase income taxes for the top 1% of the income earners (and is largely framed as such). We did not incentivize these beliefs because we were not interested in eliciting correct beliefs but in getting an estimate of their *subjective* perception of the extent of income inequality, i.e., what comes to their mind when they think about the extent of inequality.

2.3.2 Providing respondents with objective information on income inequality

We implement our information intervention after the elicitation of subjects' prior beliefs about the income inequality and several filler questions. We provide credible and objective information about the top income inequality in Switzerland to a random selection of subjects. More specifically, we randomly assign subjects to one of two groups: i) a treatment group in which subjects are provided with factual information about the degree of income inequality, or ii) a pure control group in which subjects receive no information.

Subjects in the treatment group are told the share of the total income that is received by the top 1%. We convey this information, and contrast it with their prior beliefs, using the following sentence :

"You told us that you believe that the 1% of people with the highest incomes in Switzerland receive [xx]% of the total annual income. According to the objective data collected by the Federal Department of Finance, the top 1% actually receive 12% of the total annual income."

⁶For example, if a subject first answered that the top 1% receive between 21% and 30% of the total income, they then have to provide a precise estimate within this interval.

We explicitly mention our data source (the Federal Department of Finance) in order to increase the chance that our participants believe the information we provide them with.⁷ To further illustrate the discrepancy (if any) between respondents' misperceptions and reality, we also provide them with a graphical representation that contrasts these two figures. For details on the information intervention, see Appendix A.2.

2.3.3 Measuring demand for redistribution in the context of the 99% initiative

After the information intervention, we measure all respondents' support for the 99% initiative. We start by describing the content of this initiative using a wording that is close to the wording used in the official voting booklets sent to all the Swiss voters:

“In September 2021, Switzerland will vote on the 99% initiative. The initiative aims at increasing taxation of the richest 1%. The resulting tax revenue shall be used to reduce the taxes for low and middle labor incomes or increasing social transfers. The initiative wants to reach this goal by taxing capital incomes (i.e., incomes that result from capital ownership like, for example, dividends or interest incomes from stocks and bonds) beyond a threshold at a rate that is 50% higher than labor incomes of the same amount.”

We then measure respondents' support for, or opposition to, this initiative by allowing them to make a *real monetary donation* to civic organizations that either actively support or actively oppose the 99% initiative.⁸ To that end, we endow them with 20 Swiss Francs (CHF). We first ask them whether they lean towards donating money to a civic group that supports the 99% initiative, or whether they lean towards donating to a civic group that opposes the 99% initiative. We then ask them to make a donation of up to CHF 20 to the organizations they lean towards. For example, if a subjects reveals that they lean towards donating to a civic group that opposes the 99% initiative,

⁷We measured participants' trust in the Federal Department of Finance at the end of wave 2 in order to verify that they consider the information trustworthy. Respondents could indicate how trustworthy they find the Federal Department of Finance on a scale from 1 to 7, where 1 means “not trustworthy at all” and 7 means “very trustworthy”. Overall, participants find the Swiss Federal Department of Finance very trustworthy (Mean = 5.53 ; Standard Deviation = 1.20).

⁸We purposefully did not disclose the identity of these civic groups in order to avoid that our subjects condition their donations on their subjective beliefs about the different organizations. However, we provided them with examples of such civic groups.

they have then to decide in a second stage how many—out of the CHF 20—they want to donate to such organizations. These donations are incentive compatible because the donations are effectively made to such civic groups, and because the money the subjects did not donate can be kept by themselves.⁹

2.3.4 Exit questionnaire and posterior beliefs

At the end of the second wave, we measure treated subjects' posterior beliefs about income inequality. We asked them to think about the total income that is received by all the people in Switzerland, and to indicate what they think is the share (in percent) of the total income that is received by the top 1% of people with the highest incomes. These posterior beliefs allow us to evaluate whether treated subjects update their beliefs in line with the information we provide them with. Finally, we conclude the survey with additional questions on subjects' socio-demographics and their personality traits.

2.4 Data collection, sample, and experimental protocol

Both waves were conducted online with a broadly representative sample of the German- and the French-speaking population of Switzerland and were collected in collaboration with the LINK Institute.¹⁰ In both waves, respondents were paid a show-up fee for their participation, provided that they completed the survey until the end. In the first wave, we also incentivized respondents' choices in the money allocation task by paying out each subject on the basis of their decision in one randomly chosen decision situation. In the second wave, respondents could earn whatever fraction of the CHF 20 they decided to keep instead of donating it to civic groups

⁹We decided to measure support for the 99% initiative using this two-step approach for two reasons. First, by giving subjects the possibility to make donations to organizations that *oppose* the 99% initiative, we give subjects more leeway to express their opposition to such a policy proposal, compared to just giving them the choice of only donating 0 francs to a group that supports it. Second, this two-step procedure increases the time spent thinking about the decision and thereby likely reduces the amount of random answers.

¹⁰For logistical reasons, we did not conduct the experiment in the Italian-speaking part of Switzerland, which comprises approx. 8% of the Swiss population.

campaigning in favour or against the 99% initiative.¹¹

Our final sample comprises the 1'031 subjects who participated in both waves. Descriptive statistics on participants' main socio-demographic characteristics can be found in Table A.1 in Appendix A.3. Overall, our sample is broadly representative of the Swiss voting population in the German and the French language areas with respect to age, gender, geographical area, and income. The average respondent in our sample is 47.6 years old, the share of men is 52.4%, the share of French-speaking respondents is 24.6%, and the median income is CHF 6,000. Table A.1 also shows that the control and the treatment groups are well balanced across the main observable characteristics.

An important concern in studies comprising multiple waves is selective attrition. In Table A.2 in Appendix A.4, we show that attrition between the two waves is orthogonal to the treatment assignment, to the social preference type, and to the main observable characteristics, i.e., we have very little reasons to worry about selective attrition in our sample.¹²

Another potential concern with online studies is that subjects do not pay attention to the questions they are asked and simply rush quickly through the survey. In order to measure respondents' attentiveness and to proxy data quality, we added two attention checks to the second wave. These attention checks are aimed at measuring whether participants read survey items carefully before answering them (Berinsky et al., 2014). Data quality is remarkably high in our sample: 79.2% of the subjects correctly answered both attention checks, and only 9.7% failed to pass both checks.¹³

We pre-registered our experimental design, the main hypotheses, the main outcome variables, and the sample sizes before conducting the second wave of the

¹¹Median time to complete wave 1 was 39 minutes, for which respondents were paid an average of CHF 35.5 (including a show-up fee of CHF 15; the exchange rate between points in the money allocation task and Swiss Francs was 40 points per CHF 1). Median time to complete wave 2 was 24 minutes, for which respondents were paid an average of CHF 19.7 (including a show-up fee of CHF 10).

¹²In total, 1,383 subjects participated in wave 1 in December 2020, and 1,031 subjects participated in wave 2 in May/June 2021, i.e., the rate of attrition is 25.5%.

¹³Some online samples contain considerable proportions of respondents who do not correctly answer attention check questions even in shorter surveys, with fail rates between a third and a half of the sample (see, e.g., Berinsky et al., 2014). Thus, our pass rates can be considered very high.

study.¹⁴ Ethics approval was obtained from the Human Subjects Committee of the Department of Economics of the University of Zurich (OEC IRB #2021-032).

3 The empirical distribution of social preferences

Before turning to the beliefs about inequality and their effect for demand for redistribution, we investigate the empirical distribution of social preferences in our sample. We follow Fehr *et al.* (2022, 2023) and characterize preference heterogeneity in our sample using the Dirichlet Process (DP-) means (Kulis and Jordan, 2012), a Bayesian nonparametric clustering algorithm that allows to cluster individuals into groups on the basis of their behavioral similarities.¹⁵

We apply the DP-means algorithm on the 12 distributional choices made by subjects in the money allocation task. This procedure reveals the existence of three fundamentally distinct preference types. The largest group (46.5% of the sample) comprises subjects who make predominantly payoff-equalizing choices. These subjects show both a willingness to pay to *increase* the payoff of others who are worse off, and a willingness to pay to *decrease* the payoff of others who are better off, consistent with models of inequality aversion (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). We therefore assign this cluster the label "*inequality averse*". The second largest group (38.1% of the sample) comprises individuals who display a strikingly different form of other-regarding behavior: They are also willing to pay in order to increase the payoff of those worse off, but they are generally *not* willing to pay to reduce the payoff of those who are better off, i.e., they display a high degree of altruism towards the poor but are *not* willing to reduce the income of those who are better off. This form of other-regarding behavior is consistent with an altruistic concern for the worse off (Charness and Rabin, 2002) and with altruistic other-regarding behavior that incorporates an equity-efficiency tradeoff (Fisman *et al.*, 2007, 2015). We therefore label this behavioral cluster the "*altruistic*" cluster. The last group comprises the remaining 15.4% of the individuals that make predominantly own-payoff maximizing

¹⁴The study is pre-registered on the AEA RCT registry (<https://www.socialscienceregistry.org/trials/7716>).

¹⁵We briefly summarize the key features of this approach in Appendix B.1. For an extensive discussion of this procedure and some applications, see Fehr *et al.* (2023).

choices and can therefore be labeled as being “*predominantly selfish*”. We depict and discuss the characteristic behavior of these qualitatively different behavioral types in Appendix B.2. For an extensive discussion of the identification, the characterization and the validation of these behavioral types, see Fehr [Ⓡ] al. (2022, 2023).

4 Hypotheses

Our experimental design allows us to shed light on the effects of beliefs about inequality and other-regarding preferences on demand for redistribution. It also allows us to assess whether the effects of beliefs about inequality are preference-dependent.

In previous work, Fehr [Ⓡ] al. (2022) have provided theoretical arguments and empirical evidence for the hypothesis that social preferences primarily affect individuals’ support for redistribution at higher income levels (e.g., individuals with above median incomes). Building on these results, we conjecture that correcting people’s misperceptions about the extent of inequality may primarily affect the demand for redistribution of inequality averse respondents who have an income above the median (Hypothesis 1).¹⁶ In contrast, we expect that individuals who primarily care about their own payoff as well as individuals with lower incomes will remain largely insensitive to new information about the extent of inequality. More precisely, we stipulate the following preregistered hypothesis.

HYPOTHESIS 1. A (downward) shock in beliefs about income inequality will cause a larger reduction in the demand for redistribution for inequality averse individuals with an income above the median compared to selfish individuals with an income above the median.

We preregistered this conjecture as a one-sided hypothesis because preliminary data on beliefs about the distribution of income gathered in wave 1 indicated that the majority of individuals substantially overestimate the extent of income inequality. It was therefore clear at the moment of preregistration and when we designed our

¹⁶Note also that the 99% initiative is a policy proposal that is predominantly concerned with taxing the very rich. We therefore expect that it will predominantly appeal to the inequality averse individuals who display a willingness to pay to decrease the payoff of those who are better off. We discuss the differential role of altruism and inequality aversion for the political support for different types of policy proposals in Fehr [Ⓡ] al. (2022).

information intervention that the information would shocks beliefs about income inequality of most respondents *downwards*, if anything. There is, however, no reason to expect that subjects would, on average, revise their beliefs upwards. Thus, because we are shocking beliefs about income inequality *downwards* for the vast majority of the individuals, there is no reason to expect that demand for redistribution will increase. If anything, a decrease in beliefs about inequality should either *decrease* demand for redistribution, or leave it unchanged. Because of the clear directional prediction of this conjecture, we evaluate it using one-sided tests. For all other statistical results, we use two-sided tests.

In addition, our experimental design also allows to study the extent and the ways in which beliefs and beliefs updating differ across preference types.

A relatively large literature has documented that people have biased beliefs about inequality (e.g., Cruces et al., 2013; Kuziemko et al., 2015; Karadja et al., 2017). Yet, little is known about whether these biases exist for the whole population or whether they depend on individuals' preferences. In particular, we are not aware of any empirical study that systematically relates beliefs about inequality to other-regarding preferences. While beliefs might be identically biased across the whole population, there are also good reasons to think that social preferences and beliefs about inequality are correlated. In particular, it is quite plausible that other-regarding individuals believe that there is *more* inequality than there actually is, while selfish individuals believe that there is *less* inequality than there actually is. Much like motivated beliefs, these preference-dependent beliefs might be self-sustaining: it might be easier to remain selfish by convincing oneself that there is little inequality, and it might be easier to remain other-regarding by convincing oneself that there is a lot of inequality. Thus, we conjecture that beliefs may differ by preference type. In particular, we conjecture that other-regarding individuals believe that there is more inequality than selfish individuals (Hypothesis 2).

HYPOTHESIS 2. *Other-regarding individuals overestimate the extent of inequality more than selfish individuals.*

Our last hypothesis relates to belief updating across the different preference types.

While prior beliefs might differ across preference types (Hypothesis 2), whether and how individuals update their beliefs depending on their type is unclear (Hypothesis 3). On the one hand, it is possible that the beliefs updating process depends on the preference type, and that individuals are unwilling to revise their beliefs upon seeing information that contradicts their priors. For example, other-regarding individuals might be less likely to revise their beliefs downwards if they learn that there is less inequality than they initially thought. In a similar vein, selfish individuals might be reluctant to update their beliefs upwards if they learn that there is more inequality than they initially thought. On the other hand, it is also possible that all individuals update their beliefs “rationally” upon being presented with credible and objective information about income inequality, which would suggest that the beliefs updating process is independent of preferences. However, although there is uncertainty regarding the preference-dependence and the extent of belief updating, it appears reasonable to conjecture that all preference types respond at least partly to the provision of credible information about inequality.

HYPOTHESIS 3. The information intervention corrects misperceptions about income inequality for all preference types.

5 Income inequality in Switzerland: facts and misperceptions

Before turning to the joint role of beliefs about inequality and other-regarding preferences for political support for redistribution (Hypothesis 1), we discuss the actual extent of income inequality in Switzerland. We also explore whether beliefs about inequality depend on preferences (Hypothesis 2) and the extent to which beliefs updating depends on preferences (Hypothesis 3).

5.1 The distribution of income inequality in Switzerland

Over the last century, the share of total income received by the top 1% of income earners has fluctuated around approximately 10% in Switzerland. While this share

dropped well below 10% in the late sixties, it has increased by more than 27% between 1981 and 2010 (Foellmi and Martínez, 2017). By the end of 2018, the top 1% of income earners received 12.08% of the total income in Switzerland (Swiss Federal Department of Finance, 2022). While this evolution is broadly comparable to other European countries such as Germany or France, it is in stark contrast with the United States, which have experienced a much stronger increase in the income concentration in recent years (Alvaredo et al., 2013) with the top 1% receiving close to 20% of the total US-income in 2022 (World Inequality Database, 2022).


5.2 Respondents' perceptions of income inequality

On average, respondents largely *overestimate* the share of total income received by the top 1%: They believe that the top 1% receive 54.2% of the total income (SD: 25.6 percentage points), while the actual share of 12% is almost five times lower (*t*-test that the mean equals 12%, $p < 0.001$). In terms of distribution, 93.9% of the sample overestimates inequality, while 3.8% underestimates it, and only a minority of 2.3% of the respondents has correct beliefs.

Do these beliefs depend on preferences (Hypothesis 2)? In Figure 3, we display the average prior beliefs about the income share of the top 1% as a function of respondents' preference type. The figure indicates that the beliefs are rather similar across the different preference types. On average, the inequality averse respondents believe that the share received by the top 1% is 55.1% (SD: 26.5 pp), the altruistic respondents believe that this share is 53.8% (SD: 24.2 pp), and the selfish subjects believe that it is 52.7% (SD: 27.0 pp). Thus, although the selfish individuals overestimate the income inequality slightly less than the inequality averse individuals, the differences across preference types are clearly not significant: a Kruskal–Wallis test cannot reject the null hypothesis that individuals with different preference types overestimate inequality to a similar extent ($p = 0.478$). Altogether, these results suggest that while the vast majority of our respondents vastly overestimates inequality, these misperceptions are not preference-dependent.¹⁷

¹⁷In Appendix C.1, we show that income does not predict prior beliefs. We also show that it does not predict posterior beliefs in Appendix C.2.

Figure 3: Prior beliefs about the income share of the top 1%



Note: The figure depicts the average prior beliefs about the income share of the top 1% of income-earners in Switzerland by preference type (with standard errors).

5.3 Correcting misperceptions about income inequality

Can the large and widespread misperceptions about the extent of income inequality documented above be corrected, and how do the different preference types update their beliefs (Hypothesis 3)? To assess the causal effects of beliefs about inequality, we provided all the subjects from the treatment group with credible and objective information about the share of total income received by the top 1% of income earners (see Section 2.3.2). We then elicited their knowledge about the income distribution once more at the end of the second wave. This allows us to assess whether and how respondents updated their beliefs, and whether beliefs updating depends on preference types.

Overall, the vast majority (77.3%) of treated subjects holds correct posterior beliefs, and this holds true for all preference types (for details, see Appendix C.2). In Figure 4, we depict the belief updating of the treated subjects, i.e., the difference between their posterior and their prior beliefs, as a function of their prior beliefs and their preference type. The figure shows that the vast majority of subjects who initially overestimated the income share received by the top 1% shifts their beliefs downwards by the correct magnitude, while the small share of individuals who initially under-


estimated this share shift their beliefs upwards. Importantly, individuals from all three preference types update in the correct direction and by the right magnitude. Moreover, the Figure underscores that the three preferences types update in the right direction not only on average, but for each level of prior beliefs. This is an important result, as it shows that beliefs about inequality adjust accurately, independently of preference types and irrespective of how inaccurate prior beliefs are. While one could have expected, for example, that inequality averse subjects might “want” to believe that there is a lot of inequality and be reluctant to revise their beliefs downwards even in the face of truthful information, our results show that this is clearly not the case. Overall, these results suggest that it is unlikely that different preference types hold “motivated beliefs” that help them justify their preferences.

The fact that prior beliefs and beliefs updating are independent of preferences has an important implication for our results on the effects of beliefs about inequality on support for redistribution discussed in the next section. If we find that the information intervention affects support for redistribution for only some preference types but not for others, we can rule out that this preference-specific effect is explained by differences in prior beliefs or differences in belief updating across preference types, i.e. any preference-specific effect of the information intervention on support for redistribution can *only* be explained by differences in preferences but not by differences in beliefs or beliefs updating across preference types.

6 The effects of beliefs about inequality and other-regarding preferences on support for the 99% initiative

In this section, we investigate the empirical role of social preferences and the causal effect of beliefs about income inequality for respondents’ support for the 99% initiative. We also investigate whether the effects of changes in beliefs about inequality are preference-specific.

Figure 4: Beliefs updating by prior belief and preference type



Note: This figure describes belief updating, i.e., the difference (in percentage points) between respondents' posterior beliefs and their prior beliefs, as a function of respondents' prior beliefs (x-axis) and their preference type (with standard errors). Prior beliefs are split into 5 brackets. Those who initially underestimate or precisely estimate inequality (6.1% of the subjects in the treatment group), and those who overestimate inequality (4 brackets of roughly the same size, i.e., approximately 23.5% each).

6.1 Descriptive analysis

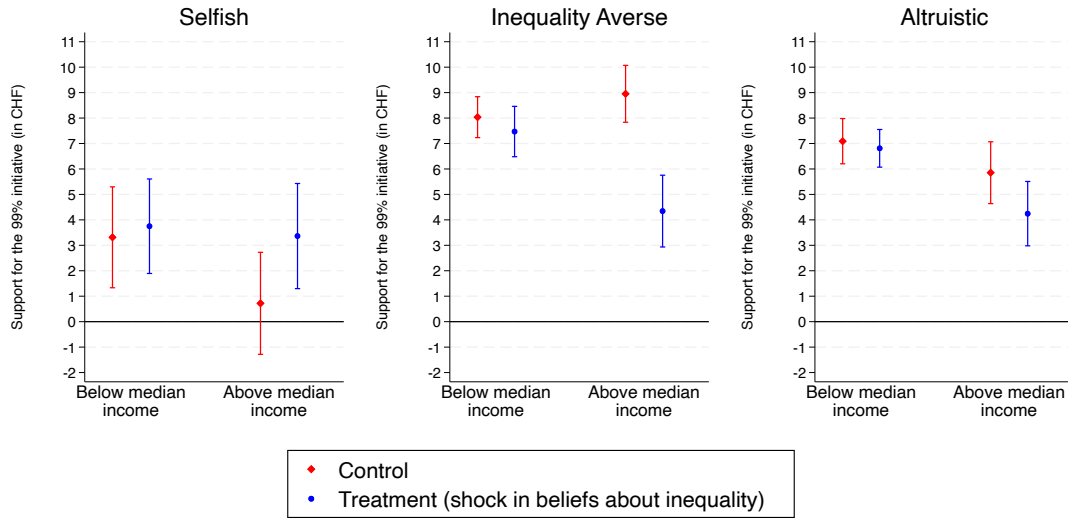
For our analysis, we code all the donations made to organizations that *oppose* the 99% initiative with negative values¹⁸ and we combine them with the donations made to organizations that *support* the 99% initiative. Our main dependent variable, which we refer to as "support for the 99% initiative", thus ranges from -20 (if the respondent makes the largest possible donation *against* the 99% initiative) to +20 (if the respondent makes the largest possible donation *in favor* of the 99% initiative). We display the distribution of this variable in Figure C.4 in Appendix C.3.

We hypothesized that a shock in beliefs would predominantly reduce the demand for redistribution of inequality averse individuals with higher incomes (Hypothesis 1). We first shed light on this question at the descriptive level. We display the average donations for the 99% initiative as a function of respondents' preference type and their income in Figure 5.

Several results are worth highlighting. First, other-regarding respondents tend

¹⁸For example, we recode a donation of CHF 15 against the 99% initiative with a value of -15.

Figure 5: The role of income and social preferences for the effect of the information shock on the donations in favor of the 99% initiative



Note: The figure shows the average donation in favor of the 99% initiative (with standard errors). Donations to organizations that oppose the 99% initiative are coded as negative values. The control group comprises subjects who were not exposed to the information intervention. The treatment group comprises subjects who were informed about the true level of income inequality.

to be a lot more supportive of the 99% initiative than selfish subjects. While selfish subjects donate an average of CHF 2.54 in favour of the 99% initiative, those with other-regarding preferences donate more than twice as much (they donate an average of CHF 6.73, test of difference: $p < 0.001$). Second, average donations decrease with income. For respondents with an income below the median, the average donation in favor of the 99% initiative is of CHF 6.89. This donation drops by more than 35%, to CHF 5.10, for respondents with an income above the median (test of difference, $p = 0.011$). Finally, the effect of the information intervention appears to be the largest amongst inequality averse subjects with an income above the median, in line with Hypothesis 1. While inequality averse respondents with an income above the median donate an average of CHF 8.95 in the control group, those in the treatment group donate an average of CHF 4.34 (test of difference, $p = 0.012$), i.e., the affluent inequality averse subjects who revise their beliefs about inequality downwards display much lower support for the 99% initiative than those whose beliefs are not shocked downwards. This treatment effect is much larger than the treatments effects documented amongst individuals with an income below the median, or amongst other

social preference types, consistent with Hypothesis 1.

6.2 Regression analysis

To shed further light on the causal effect of beliefs about inequality and their interaction with other-regarding preferences, we estimate the following model

$$\begin{aligned} \text{Support}_i = & \beta_0 + \beta_1 \text{Treatment}_i + \beta_2 \text{IA}_i + \beta_3 \text{Altruistic}_i + \\ & + \beta_4 \text{Treatment}_i \times \text{IA}_i + \beta_5 \text{Treatment}_i \times \text{Altruistic}_i + \Gamma' X_i + \varepsilon_i \end{aligned} \quad (1)$$

where Support_i is our measure of support for the 99% initiative based on subjects' donations. Treatment_i is an indicator variable that takes the value one if the respondent is in the information treatment. IA_i is a dummy that takes the value one if the respondent is inequality averse, and Altruistic_i is a dummy that takes the value one if the respondent is altruistic. The two interaction variables, $\text{Treatment}_i \times \text{IA}_i$ and $\text{Treatment}_i \times \text{Altruistic}_i$, are aimed at capturing the possible interactions between the treatment and the two other-regarding types. The omitted category in these regressions are the predominantly selfish individuals assigned to the control group. For some of our regressions, we also include a rich set of individual-level controls, X_i , which comprise respondents' socio-demographics, proxies for their economic preferences and general trust (Falk et al., 2022), proxies for their financial situation, as well as their prior beliefs about the income share of the top 1%, the determinants of success, the income and wealth distributions in Switzerland, financial mobility, poverty in Switzerland, and distrust in politicians. Finally, ε_i is an individual-specific error term.

We report the results of our estimates in Table 1.¹⁹ Columns 1 and 2 show the average effect of the information intervention on the full sample. On average, a downward shock in beliefs about inequality causes a small but insignificant reduction in donations in favor of the 99% initiative ($p = 0.106$). This result is consistent with a recent meta-analysis showing that while presenting subjects with information about inequality generally yields large changes in beliefs and successfully corrects misper-

¹⁹For transparency, we also depict the full regression results in Table C.1 in Appendix C.4.

ceptions, it very often does *not* substantially affect demand for redistribution (Ciani et al., 2021).

This result hides, however, substantial heterogeneity. For example, column 3 shows that the information shock has a large negative effect on the average donations of inequality averse subjects (CHF -4.30, approximately -0.4 of a standard deviation, $p = 0.045$), but that it does not affect the donations of the selfish and the altruistic subjects. In addition, the results also indicate that inequality averse respondents donate significantly more to support the 99% initiative than selfish subjects (CHF +6.87 or +0.62 of a standard deviation, $p < 0.001$), and so do altruistic subjects (CHF +5.00 or +0.45 of a standard deviation, $p = 0.001$). Column 4 shows that these effects survive the inclusion of a large set of control variables.

What is the joint role of beliefs about inequality, preferences and income for support for redistribution? Do beliefs about inequality predominantly affect the demand for redistribution of the inequality averse respondents who have higher incomes? We conjectured (Hypothesis 1) that the information intervention will predominantly reduce the demand for redistribution of inequality averse individuals with above-median incomes, relative to selfish individuals with above-median incomes. To test this hypothesis, we estimate equation (1) separately for respondents with an income below the median, and for respondents with an income above the median. We depict the results of these estimations in columns 5-8.²⁰

We find that the association between other-regarding preferences and support for the 99% initiative is rather small and insignificant among individuals with an income below the median (see, e.g., column 6), but large in magnitude and strongly significant among those with an income above the median (columns 7-8).²¹

Turning to the effect of the information intervention, we find that our treatment does *not* affect the demand for redistribution of subjects with an income below the

²⁰51 subjects did not disclose their income and are thus not included in columns 5-8.

²¹The inequality averse subjects with an above-median income donate significantly more than the selfish (CHF +6.41 or 0.58 of a standard deviation, $p = 0.003$). Likewise, the altruistic subjects with an above-median income donate significantly more than the selfish (CHF +5.43 or 0.49 of a standard deviation, $p = 0.014$). In contrast, the coefficient for altruistic subjects with an income below the median is insignificant ($p = 0.140$) and the one for inequality averse subjects with an income below the median is only very weakly significant ($p = 0.067$).

median (regardless of their preference type), but that it significantly affects the demand for redistribution of the inequality averse subjects with an income above the median. As conjectured in Hypothesis 1, for these individuals a downward shock in beliefs about inequality leads to a large and strongly significant decrease in support for the 99% initiative (columns 7-8, CHF -6.54 or -0.59 of a standard deviation, $p = 0.022$ for $H_0 : \beta_4 \geq 0$, see the p-values for the relevant test at the bottom of the Table). In contrast, the information intervention does *not* affect the donations of the selfish and the altruistic subjects with an income above the median.

Overall, these results highlight the heterogeneous effects of beliefs about inequality and other-regarding preferences for demand for redistribution. While all subjects tend to have biased beliefs about the extent of income inequality, correcting these misperceptions only significantly affects the demand for redistribution of the affluent inequality averse respondents.

6.3 Robustness: Attention and experimenter demand effects

One potential concern with information experiments conducted online is that some individuals might not pay attention to the information that is presented to them, which could considerably alter the results of the study (Berinsky et al., 2014). To account for this potential problem, our study included attention checks. Less than 10% of the subjects failed to correctly answer any of our screener questions, which is very reassuring and relatively low compared to related studies.²² Overall, our main results are robust to excluding subjects who did not pass the attention checks (see Appendix C.5).

Another potential concern with information provision studies are experimenter demand effects (see, e.g., Haaland et al., 2023). While demand effects can be a serious concern in some settings, we believe that they are unlikely to be a threat to our results due to the preference-specificity of the predictions we are testing. Indeed, we hypothesized (Hypothesis 1) that a shock in beliefs about inequality would predominantly decrease the demand for redistribution of inequality averse subjects with an

²²For example, between 30 and 50 percent of the participants fail to pass the attention checks in the studies reported in Berinsky et al. (2014).

Table 1: Determinants of donations in favor of the 99% initiative

	Full sample				Below median income		Above median income	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	-1.106 (0.684)	-0.999 (0.668)	1.983 (1.881)	2.042 (1.825)	0.436 (2.694)	1.727 (2.815)	2.643 (2.860)	2.553 (2.718)
IA			6.871*** (1.519)	5.341*** (1.514)	4.721** (2.123)	4.501* (2.454)	8.232*** (2.288)	6.407*** (2.130)
Altruistic			5.001*** (1.554)	4.350*** (1.536)	3.777* (2.156)	3.730 (2.523)	5.133** (2.337)	5.434** (2.194)
Treat x IA			-4.298** (2.145)	-3.909* (2.091)	-1.001 (2.981)	-2.315 (3.077)	-7.252** (3.381)	-6.542** (3.247)
Treat x Altruistic			-2.951 (2.129)	-3.258 (2.066)	-0.716 (2.932)	-2.289 (3.069)	-4.252 (3.356)	-4.195 (3.266)
Constant	6.644*** (0.471)	-0.708 (4.592)	1.568 (1.376)	-4.795 (4.875)	3.314* (1.964)	-1.754 (6.243)	0.721 (1.996)	-11.661 (9.310)
Beliefs	No	Yes	No	Yes	No	Yes	No	Yes
Socio-demographics	No	Yes	No	Yes	No	Yes	No	Yes
Education	No	Yes	No	Yes	No	Yes	No	Yes
Occupation	No	Yes	No	Yes	No	Yes	No	Yes
Preference measures	No	Yes	No	Yes	No	Yes	No	Yes
Wealth bracket dummies & financial assets	No	Yes	No	Yes	No	Yes	No	Yes
p-value(Ho: Treat x IA \geq 0)			0.023	0.031	0.369	0.226	0.016	0.022
R ²	0.003	0.132	0.029	0.146	0.020	0.136	0.038	0.232
Observations	1031	1030	1031	1030	558	557	422	422

Note: OLS regression. The dependent variable is the donation amount towards an organization in favor of the 99% initiative for the full sample (columns 1-4), for subjects with an income below the median (columns 5-6), and for subjects with an income above the median (columns 7-8). Subjects who did not disclose their income are not included in columns 5-8. Beliefs include subjects' prior beliefs about the income share of the top 1%, their prior beliefs about income and wealth distributions in Switzerland, as well as their beliefs about the determinants of success, financial mobility, poverty in Switzerland, and distrust in politicians. Other socio-demographics include age, age squared, a dummy variable indicating whether the respondent is male, a dummy variable indicating whether the respondent's native language is French, and a dummy indicating whether the respondent is married. Education includes dummies indicating a respondent's highest educational achievement (compulsory school, vocational training, high school, university, or other). Occupation includes dummies indicating a respondent's occupation status (currently has a full-time job, a part-time job, is a student, is a pensioner, is unemployed, or other), a dummy variable indicating whether the respondent has experienced unemployment in the past. Preference measures from the global preference survey (Falk et al., 2022) include measures of risk aversion, (positive and negative) reciprocity, impatience, and general trust in strangers. Wealth bracket dummies and financial assets include dummy variables for each wealth level, and a dummy variable indicating whether the respondent has invested part of her wealth in financial assets. For a detailed explanation of the covariates, see Appendix A.1. Robust standard errors are reported in parentheses. p-value(Ho: Treat \times IA \geq 0) indicates the p-value associated with a one-sided test evaluating the null hypothesis that the information intervention has a non-negative effect on the donations of the inequality averse subjects. Levels of significance: * $p < .1$, ** $p < .05$, *** $p < .01$.

income above the median. If our results were predominantly driven by experimenter demand effects, we should observe that *all* subjects adjust their demand for redistribution as a response to our information treatment, irrespective of their preference type and their income. This is however *not* what we find. Instead, the information intervention only affects the demand for redistribution of the affluent inequality averse subjects—consistent with what we conjectured. In addition, recent methodological contributions have shown that strong demand effects (generated on purpose by the experimenter) result in only very modest behavioral responses in similar survey experiments (de Quidt et al., 2018).

7 Concluding remarks

Over the last decade, several studies have highlighted the role of beliefs about inequality as well as the role of other-regarding preferences for support for redistribution. These two strands of the literature have largely evolved separately. In this paper, we studied the nature of support for redistribution by exploring the *joint* role of social preferences and beliefs about inequality. We also explored whether beliefs about inequality and beliefs updating depend on preferences. We investigated these questions by conducting an online experiment with a representative sample of the Swiss population in the context of the 99% initiative, a highly redistributive policy proposal aimed at increasing taxes on the richest individuals in Switzerland.

We showed that the vast majority of individuals largely overestimate the extent of income inequality in Switzerland, and that these misperceptions are independent of preference types. We also showed that all subjects update their beliefs correctly upon receiving information about the true share of total income received by the top 1% of income earners, irrespective of their preferences.

Turning to the effects of beliefs about inequality for support for redistribution, we showed that correcting biased beliefs about inequality has a quantitatively large and significant effect on the support for the 99% initiative for inequality averse individuals with an income above the median, but that it essentially does *not* affect other respondents—consistent with our pre-registered hypothesis. These results highlight

the joint importance of beliefs about inequality and other-regarding preferences for demand for redistribution. They also underscore the possibly large heterogeneous effects that information interventions might generate and that remain hidden in aggregate analyses. In particular, these heterogeneous effects could explain the somewhat puzzling finding that information interventions often successfully generate large corrections in beliefs about inequality *without* substantially affecting demand for redistribution, as documented in a recent meta-analysis (Ciani et al., 2021).

References

- Alesina, Alberto and George-Marios Angeletos (2005) "Fairness and redistribution," *American Economic Review*, 95 (4), 960–980.
- Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso (2018) "Intergenerational mobility and preferences for redistribution," *American Economic Review*, 108 (2), 521–54.
- Almås, Ingvild, Alexander W Cappelen, and Bertil Tungodden (2019) "Cutthroat capitalism versus cuddly socialism: Are Americans more meritocratic and efficiency-seeking than Scandinavians?" *NHH Dept. of Economics Discussion Paper*.
- Alvaredo, Facundo, Anthony B Atkinson, Thomas Piketty, and Emmanuel Saez (2013) "The top 1 percent in international and historical perspective," *Journal of Economic perspectives*, 27 (3), 3–20.
- Bellemare, Charles, Sabine Kröger, and Arthur Van Soest (2008) "Measuring inequity aversion in a heterogeneous population using experimental decisions and subjective probabilities," *Econometrica*, 76 (4), 815–839.
- Benabou, Roland and Efe A Ok (2001) "Social mobility and the demand for redistribution: the POUM hypothesis," *The Quarterly Journal of Economics*, 116 (2), 447–487.
- Benabou, Roland and Jean Tirole (2006) "Belief in a just world and redistributive politics," *The Quarterly Journal of Economics*, 121 (2), 699–746.
- Berinsky, Adam J, Michele F Margolis, and Michael W Sances (2014) "Separating the shirkers from the workers? Making sure respondents pay attention on self-administered surveys," *American Journal of Political Science*, 58 (3), 739–753.
- Bolton, Gary E and Axel Ockenfels (2000) "ERC: A theory of equity, reciprocity, and competition," *American Economic Review*, 90 (1), 166–193.
- Bruhin, Adrian, Ernst Fehr, and Daniel Schunk (2018) "The many faces of human sociality: Uncovering the distribution and stability of social preferences," *Journal of the European Economic Association*, 17 (4), 1025–1069.

- Charness, Gary and Matthew Rabin (2002) "Understanding social preferences with simple tests," *The Quarterly Journal of Economics*, 117 (3), 817–869.
- Ciani, Emanuele, Louis Fréget, and Thomas Manfredi (2021) "Learning about inequality and demand for redistribution: A meta-analysis of in-survey informational experiments."
- Comiter, Marcus Zachary, Miriam Cha, HT Kung, and Surat Teerapittayanon (2016) "Lambda means clustering: automatic parameter search and distributed computing implementation."
- 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.
- Falk, Armin, Anke Becker, Thomas J Dohmen, David Huffman, and Uwe Sunde (2022) "The Preference Survey Module: A Validated Instrument for Measuring Time, Risk, and Social Preferences," *Management Science*, 69 (4), 1935–1950.
- Fehr, Dietmar, Johanna Mollerstrom, and Ricardo Perez-Truglia (2019) *Your place in the world: The demand for national and global redistribution*: National Bureau of Economic Research.
- Fehr, Ernst & Thomas Epper & Julien Senn (2022) "Other-regarding preferences and redistributive politics," 10.5167/UZH-183052.
- (2023) "The fundamental properties, stability and predictive power of distributional preferences."
- Fehr, Ernst and Klaus M Schmidt (1999) "A theory of fairness, competition, and cooperation," *The Quarterly Journal of Economics*, 114 (3), 817–868.
- Fisman, Raymond, Pamela Jakiela, and Shachar Kariv (2017) "Distributional preferences and political behavior," *Journal of Public Economics*, 155, 1–10.
- Fisman, Raymond, Pamela Jakiela, Shachar Kariv, and Daniel Markovits (2015) "The distributional preferences of an elite," *Science*, 349 (6254), aab0096.

- Fisman, Raymond, Shachar Kariv, and Daniel Markovits (2007) "Individual preferences for giving," *American Economic Review*, 97 (5), 1858–1876.
- Foellmi, Reto and Isabel Z Martínez (2017) "Volatile top income shares in Switzerland? Reassessing the evolution between 1981 and 2010," *Review of Economics and Statistics*, 99 (5), 793–809.
- Fong, Christina (2001) "Social preferences, self-interest, and the demand for redistribution," *Journal of Public Economics*, 82 (2), 225–246.
- Haaland, Ingar and Christopher Roth (2020) "Labor market concerns and support for immigration," *Journal of Public Economics*, 191, 104256.
- Haaland, Ingar, Christopher Roth, and Johannes Wohlfart (2023) "Designing information provision experiments," *Journal of economic literature*, 61 (1), 3–40.
- 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), 201–212.
- Kerschbamer, Rudolf and D Müller (2020) "Social preferences and political attitudes: An online experiment on a large heterogeneous sample," *Journal of Public Economics*, 182.
- Kulis, Brian and Michael I. Jordan (2012) "Revisiting k-means: New Algorithms via Bayesian Nonparametrics," *Proceedings of the 29th International Conference of Machine Learning*.
- 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.
- Lergetporer, Philipp, Guido Schwerdt, Katharina Werner, Martin R West, and Ludger Woessmann (2018) "How information affects support for education spending: Evidence from survey experiments in Germany and the United States," *Journal of Public Economics*, 167, 138–157.

- Piketty, Thomas (1995) "Social mobility and redistributive politics," *The Quarterly Journal of Economics*, 110 (3), 551–584.
- de Quidt, Jonathan, Johannes Haushofer, and Christopher Roth (2018) "Measuring and bounding experimenter demand," *American Economic Review*, 108 (11), 3266–3302.
- Roth, Christopher, Sonja Settele, and Johannes Wohlfart (2021) "Beliefs about public debt and the demand for government spending," *Journal of Econometrics*, 231 (1), 165–187.
- Settele, Sonja (2022) "How do beliefs about the gender wage gap affect the demand for public policy?" *American Economic Journal: Economic Policy*, 14 (2), 475–508.
- World Inequality Database (2023) "World Inequality Database," <https://wid.world/>, Accessed: 2023-11-28 17:44.

A Background information on experimental tasks, survey measures and population sample

A.1 Details on the measurement of covariates from wave 1

Socio-demographics

We collected information on age, gender, language region, marital status, the highest achieved level of education (compulsory school, vocational training, high school, university, other), occupational status (full time job, part-time job, student, pensioner, currently unemployed, other), whether the individual has experienced unemployment in the past, and municipality of residence.

Belief about own control over success

“People differ in their views regarding why some people get ahead and succeed in life while others do not succeed. Please tell us how important you think each of the factors listed below is for why some people get ahead and succeed in life. For each factor, please give your answer on a scale from 0 to 10, where 0 means “not at all important” and 10 means “extremely important”. You can choose any number between 0 and 10.” (Source: Fong (2001) and Gallup)

- Willingness to take risks
- Inherited wealth (reverse coded)
- Hard work and initiative
- Luck, being at the right time at the right place (reverse coded)
- Striving for the right education and training


We then create an index by averaging out these five items.

Belief about intergenerational income mobility

“We would like to ask you what you think about the life opportunities of children born in the poorest families in Switzerland. For the following question, we focus on 500 families

that represent all the Swiss families with children. We divide them into five groups on the basis of their income, with each group containing 100 families. These groups are: the richest 100 families, the second richest 100 families, the middle 100 families, the second poorest 100 families, and the poorest 100 families. All these groups are depicted in the figure below, ranked from the richest families to the poorest families. In your opinion, out of 100 children coming from the poorest 100 families, how many will belong to each of the five income groups depicted in the picture below once they become adults? Please fill out the entries to the right of the figure below. Note that your entries need to add up to 100 and that no decimals are allowed. From our experience, this question takes a bit of time to be answered carefully.”(Source: Alesina et al. (2018))

Figure A.1: Question to elicit belief of intergenerational income mobility



Belief about degree of poverty in Switzerland








“According to the Swiss Federal Office for Statistics, the poverty line in Switzerland is equal to 2’293 francs per month for a single-person household, and 3’968 francs per month for a household with two adults and two children. This means that any single-person household living with less than 2’293 francs per month is considered as “living in poverty”. Similarly, a person living in a household consisting of 2 children and 2 adults with an income of less than 3’968 francs per month is considered as “living in poverty”. Currently, Switzerland has a population of 8.6 million people. In your opinion, how many people in Switzerland currently

live with an income below the poverty line?" [1. Less than 100'000 people, 2. Between 100'001 and 200'000 people, 3. Between 200'001 and 300'000 people, 4. Between 300'001 and 400'000 people, 5. Between 400'001 and 500'000 people, 6. Between 500'001 and 600'000 people, 7. Between 600'001 and 700'000 people, 8. Between 700'001 and 800'000 people, 9. Between 800'001 and 900'000 people, 10. Between 900'001 and 1 million people, 11. Above 1 million people]

Belief about income distribution

"The next question refers to your perception of the income received by different groups of people in Switzerland. By "income" we mean all the revenue people receive for their work, but also the revenue they get from their investments (e.g., returns on bonds, stocks or bank account) or what they receive from the state (e.g., pensions or other welfare benefits). For this question, we focus on 100 individuals that represent the Swiss population. We divide these 100 people into 7 different groups, ranging from the 1% of the people with the highest income in Switzerland (the top 1%), to the 20% of the people with the lowest income in Switzerland (the bottom 20%). All these groups are depicted in the figure below, ranked from the 1% with the highest income to the 20% with the lowest incomes. Think about the total income that is received by all the people in Switzerland. In your opinion, what percent of the total income in Switzerland does each of the groups shown in the figure below receive? Please fill out the entries to the right of the figure below. Note that your entries need to add up to 100 and that no decimals are allowed. From our experience, this question takes a bit of time to be answered carefully."

Figure A.2: Question to elicit belief of income distribution

	The 1% highest-earning people		receive the following share (%) of the total income: <input type="text"/>
	The following 9% (top 2-10%) of income earners		receive the following share (%) of the total income: <input type="text"/>
	The following 10% (top 11-20%) of income earners		receive the following share (%) of the total income: <input type="text"/>
The following 20% (top 21-40%) of income earners			receive the following share (%) of the total income: <input type="text"/>
The middle 20% of income earners			receive the following share (%) of the total income: <input type="text"/>
The second lowest 20% of income earners			receive the following share (%) of the total income: <input type="text"/>
The 20% lowest-earning people			receive the following share (%) of the total income: <input type="text"/>
			TOTAL <input type="text"/>

Belief about wealth distribution

“The next question refers to your perception of how much wealth is possessed by different groups of people in Switzerland. By “wealth”, also known as net worth, we mean the total value of everything someone owns minus any debt that he or she owes. A person’s net worth includes the sum of all their savings and all their other assets such as real estate property, stocks, bonds, art collections, etc., minus the sum of all their liabilities such as loans and mortgages. For this question, we focus on 100 people that represent the Swiss population. We divide these 100 people into 7 different groups, ranging from the 1% of the people with the highest wealth in Switzerland (the richest 1%) to the 20% of the people with the lowest wealth in Switzerland (the poorest 20%). All these groups are depicted in the figure below, ranked from the 1% with the highest wealth to the 20% with the lowest wealth. Think about all the wealth that is possessed by all the people in Switzerland. In your opinion, what percent of total wealth in Switzerland is possessed by each of the groups shown in the figure below? Please fill out the entries to the right of the figure below. Note that your entries need to add up to 100 and that no decimals are allowed. From our experience, this question takes a bit of time to be answered carefully.”

Figure A.3: Question to elicit belief of wealth distribution

the richest 1% } possess the following share (%) of the total wealth:

the richest 2 to 10% } possess the following share (%) of the total wealth:

the richest 11 to 20% } possess the following share (%) of the total wealth:

the second richest 20% } possess the following share (%) of the total wealth:

the middle 20% } possess the following share (%) of the total wealth:

the second poorest 20% } possess the following share (%) of the total wealth:

the poorest 20% } possess the following share (%) of the total wealth:

TOTAL

Distrust in politicians

"What do you think about the following statement? Swiss politicians work to enrich themselves and the lobbies that they support instead of working for the benefit of the majority of the citizens." [1. Completely disagree, 2. Disagree, 3. Rather disagree, 4. Neither agree nor disagree, 5. Rather Agree, 6. Agree, 7. Absolutely agree]

Preference measures and trust

We measured risk preferences, patience, negative reciprocity and positive reciprocity, as well as subjects' general trust in people with the experimentally validated survey questions of Falk et al. (2022).

Financial situation

Own income. *"We now turn to a few questions that relate to your income and your current financial situation. By "income" we mean all the revenue you receive for your work, but also the revenue you get from your investments (e.g., returns on bonds, stocks or bank account) or what you receive from the state (e.g., pensions or other welfare benefits). How much was your income last month (before taxes)? [Less than 3'000 francs, between 3,001 and 4,000 francs,*

between 4,001 and 5,000 francs, . . . , between 13'001 and 14'000 francs, between 14'001 and 15'000 francs, more than 15'000 francs, No answer]

Income mobility Two mobility measures are constructed from the following three questions (based on Fong (2001) and Gallup):

1. *Think of the picture of a ladder. Suppose we say that step 10, which is at the top of the ladder, represents the group with the highest income in Switzerland, and that step 0, at the bottom of the ladder, represents the group with the lowest income in Switzerland. On which step of the ladder do you feel you personally stand at the present time? [0, . . . , 10] [current step]*
2. *On which step would you say you stood five years ago? [0, . . . , 10] [past step]*
3. *Just your best guess, on which step do you think you will stand in five years? [0, . . . , 10] [future step]*

Based on these questions, we create the following two measures:

- Beliefs about future mobility = future step - current step.
- Perceived past mobility = current step – past step.

Own wealth *"We now turn to a few questions that relate to your wealth and the wealth you might have inherited from your parents. By "wealth", also known as net worth, we mean the total value of everything you own minus any debt that you owe. Your net worth includes the sum of all your savings and all your other assets such as real estate property, stocks, bonds, art collections, etc., minus the sum of all your liabilities such as loans and mortgages. Think about the sum of everything you own, minus the debt you owe. How much do you estimate is your current net worth? [My debt exceeds what I own, Between 0 and 25'000 francs, Between 25'001 and 50'000 francs, Between 50'001 and 75'000 francs, Between 75'001 and 100'000 francs, Between 100'001 and 200'000 francs, Between 200'001 and 500'000 francs, Between 500'001 and 1'000'000 francs, Between 1'000'001 and 2'000'000 francs, more than 2'000'000 francs, No answer]*


Capital investments *Is part of your wealth invested in funds, shares, bonds, and similar financial assets? (For this question, please ignore retirement provisions that relate to your 2. pillar.) [Yes,No]*

A.2 Information intervention

The information intervention consists of two separate screens. On the first screen, treated subjects receive the following introductory message: *“In recent years, questions related to the distribution of income have been frequently discussed in the society. When discussing such issues, it is important to have accurate information. The Federal Department of Finance collects data that provide an objective measure of the extent of income inequality in Switzerland. As you might not be aware of these numbers, we will reveal them to you in the next screen.”* The information is provided on the second screen (see the screenshot below). It contains four elements: 1. We tell subjects whether they are [overestimating, underestimating, correctly estimating] the income share of the top 1%; 2. We remind them of their prior belief; 3. We inform them about the true share of total income that is received by the top 1% of people with the highest incomes in Switzerland using the latest objective data collected by the Federal Department of Finance; 4. We also provide a graphical representation of how their prior belief compares with the truth.

Figure A.4: Presentation format of the information intervention (example of a participant who overestimates inequality)

You are overestimating the income share that the top 1% of people with the highest incomes in Switzerland receive. You told us that you believe that they receive **58%** of the total annual income (**red bar** in graph below). According to the objective data collected by the Federal Department of Finance, the top 1% actually receive **12%** of the total annual income (**green bar** in graph below).



A.3 Demographic characteristics of sample population

We depict the main descriptive statistics in Table A.1, separately for the treatment and the control group. The last column indicates that our treatment is well balanced across the main observable characteristics, as well as across preference types. The table also indicates that our sample is broadly representative of the Swiss population with respect to age, gender, geographical area, and income.

Table A.1: Descriptive statistics and balance checks

	Population	Treatment	Control	<i>p</i> -value (<i>t</i> -test)
Age (mean)	51.1	46.8	48.3	.101
Male	0.48	0.53	0.52	.702
French-speaking	0.25	0.27	0.23	.143
Income bracket : \leq CHF 4000	0.28	0.32	0.32	.852
Income bracket : CHF 4001-6000	0.26	0.22	0.22	.841
Income bracket : CHF 6001-8000	0.22	0.21	0.19	.490
Income bracket : CHF 8001-10000	0.12	0.11	0.11	.978
Income bracket : CHF 10001-15000	0.09	0.06	0.08	.173
Income bracket : \geq CHF 15000	0.03	0.02	0.03	.628
Income bracket : NA	-	0.05	0.05	.541
Above-median belief about own control over success	-	0.46	0.46	.957
Prior belief about income share of top 1%	-	52.9	55.6	.091
Above-median prior belief about income share of top 1%	-	0.47	0.52	.113
Inequality Averse type	-	0.46	0.47	.902
Altruistic type	-	0.39	0.37	.641
Selfish type	-	0.15	0.16	.647

Note: The table displays descriptive statistics of the Swiss population and of our sample, separately for the treatment and the control group. The descriptive statistics include age (mean), the share of male people, the share of French-speaking people, as well as the shares of people falling into each monthly income bracket. The population data were obtained from the Swiss Federal Bureau of Statistics (2018) and are restricted to the adult Swiss population (i.e., individuals holding a Swiss passport who are at least 18 years old). In addition, the descriptive statistics include the share of subjects with above-median beliefs about own control over success in life, the prior belief about the income share of the top 1% (percentage), as well as the share of subjects with above-median prior beliefs about the income share of the top 1%. Finally, the descriptive statistics include the shares of inequality averse, altruistic, and selfish subjects.

A.4 Attrition

In the Table A.2, we show that participation in the second wave is orthogonal to the treatment, to the preference type, and to the bulk of the observable characteristics. Note that older respondents are slightly more likely *not* to drop out between waves.

Table A.2: Participation in 2021 wave

	(1)	(2)	(3)	(4)	(5)	(6)
Treatment	0.016 (0.023)					0.020 (0.023)
IA		0.022 (0.034)				0.024 (0.035)
Altruistic		-0.030 (0.036)				0.015 (0.036)
Income bracket: \leq CHF 4000			-0.014 (0.057)			0.001 (0.057)
Income bracket: CHF 4001-6000			0.026 (0.059)			0.017 (0.058)
Income bracket: CHF 6001-8000			0.070 (0.059)			0.039 (0.059)
Income bracket: CHF 8001-10000			0.009 (0.064)			-0.037 (0.064)
Income bracket: CHF 10001-15000			0.041 (0.069)			-0.008 (0.070)
Income bracket: $>$ CHF 15000			-0.043 (0.095)			-0.090 (0.094)
Above-median belief of own control over success				0.009 (0.024)		0.007 (0.023)
Male					0.029 (0.023)	0.035 (0.025)
Age					0.005*** (0.001)	0.005*** (0.001)
French-speaking					-0.020 (0.027)	-0.020 (0.027)
Constant	0.737*** (0.017)	0.748*** (0.030)	0.729*** (0.053)	0.741*** (0.016)	0.509*** (0.041)	0.468*** (0.075)
R^2	0.000	0.003	0.005	0.000	0.033	0.038
Observations	1383	1383	1383	1383	1383	1383

Note: OLS regression. The dependent variable equals 1 if the individual from the first wave participated in the second wave and equals 0 if the individual did not participate. Treatment is a dummy variable indicating whether the respondent is randomized into the information treatment. IA is a dummy that takes the value 1 if the respondent is inequality averse, and Altruistic is a dummy that takes the value 1 if the respondent is altruistic. Income brackets are dummy variables that equal 1 if the respondent falls into the respective monthly income category. Above-median beliefs about own control over success is a dummy variable indicating whether a subject has high beliefs of own control over success in life. Further socio-demographic variables include a dummy variable indicating whether the respondent is male, age, and a dummy variable indicating whether the respondent's native language is French. Levels of significance: * $p < .1$, ** $p < .05$, *** $p < .01$.

B Identifying preference types using DP-Means

B.1 The method

We identify heterogeneity in preferences by applying a nonparametric Bayesian approach—the Dirichlet Process (DP) means clustering algorithm (Kulis and Jordan, 2012). This appendix only provides a brief overview of this method. For a more detailed description of the DPM algorithm and for a discussion of its key differences with other clustering methods, see Fehr ^(r) al. (2022, 2023).

This DP-means algorithm groups individuals into clusters according to their *behavioral similarities*. In our context, clusters are based on subjects’ 12 distributional choices in the money allocation task (Figure 2a), and similarity is measured by “how close” an individual’s allocation profile is to the average allocation of a cluster. Our implementation of the algorithm is based on an iterative refinement. We first span an m -dimensional space, with m denoting the number of budget lines used for the clustering algorithm (in our case, $m = 12$, the twelve budget lines presented in Figure 2a). Each individual’s choices are represented by a single point in the 12-dimensional space. We then ask how subjects populate this space. Specifically, we are interested in the number of clusters (i.e. types) that emerge and individuals’ assignment to clusters. A cluster is characterized by the set of the individuals assigned to the cluster and the associated mean vector of observations (the “centroid”), which – in our case – represents the mean (cluster- representative) behavior of all individuals in m -dimensional space that belong to the cluster.

We initialize the algorithm with a single centroid specified as the global mean vector. At this stage, all data points are assigned to this single centroid. We then refine by iterating over the following two steps: First, we sequentially go through the list of data points in m -dimensional space (i.e. subjects), and check for each subject whether any of the squared Euclidean distances to the centroid exceeds a cluster penalty parameter. If this is the case, we open up a new cluster with the actual data point’s location vector as the centroid. Otherwise, we assign the data point to its nearest cluster. Second, we collect the subjects assigned to the same clusters and update the centroids by computing the mean vector for each cluster. These two steps

are repeated until convergence is reached, i.e. until there is no change in subjects' assignments.

An important aspect of the DP-means approach is that it enables the identification of preference types without committing to a prespecified number of different preference types. Moreover, this approach does neither require an ex-ante specification or parameterization of types, nor does it presume a specific error structure. This means that it remains ex-ante agnostic about key distributional assumptions, and it does not constrain heterogeneity to lie within a predetermined set of models or parameter space.²³ The DP-means algorithm allows for all possible type partitions of the data spanning from a representative agent up to as many types as there are individuals in the population, i.e., it determines the number of preferences types endogenously. Thus, (i) the actual number of types, (ii) the assignment of each individual to one of the types and (iii) the behavioral (preference) properties of the types emerge endogenously.²⁴

²³In this regard, our approach differs from previous work (e.g. Bellemare et al., 2008; Fisman et al., 2015, 2017; Bruhin et al., 2018) that characterized preference heterogeneity on the basis of structural assumptions on preferences and error terms.

²⁴The fact that the number of types adapts to the data has important benefits (see Kulis and Jordan, 2012). Most notably, as previous work has shown (see Comiter et al., 2016), this feature of the algorithm yields higher quality type-separation than methods that specify the number of types prior to clustering (such as *k*-means).

B.2 Distribution of choices for each preference type

The application of the DP-means algorithm to the money allocation task in our general population sample suggests the existence of *three* behavioral types. Roughly half of the subjects (46.5%) are assigned to Type 1, around one-third (38.1%) to Type 2, and the remainder (15.4%) to Type 3. The three types differ substantially in terms of their behavior. A careful examination of the decisions of these types permits us to assign them a label with a clear behavioral interpretation.

Figure B.1 depicts the relative share of own-payoff minimizing, payoff-equalizing, and own-payoff-maximizing choices, for each identified cluster.²⁵ Positively sloped budget lines inform us on subjects' willingness to pay to decrease the payoff of those better off, while negatively sloped budget lines inform us on their willingness to pay to increase the payoff of those worse off.

The figure shows that individuals in type 1 predominantly make payoff-equalizing choices, both for negatively sloped budget lines and for positively sloped budget lines. They thus exhibit a willingness to pay (i) for reducing inequality when this involves increasing the other individual's payoff (i.e., for negative slopes) and (ii) when it involves decreasing the other individual's payoff (i.e., for positive slopes). For this reason, we assign them the label *inequality averse*.

This pattern contrasts sharply with individuals assigned to type 2, who display a substantial willingness to pay to increase the other individual's payoff (negatively sloped budget lines), but are generally unwilling to pay to decrease the other's payoff (budget lines with positive slope). We therefore label individuals in this cluster as "altruists."


Last, individuals in the third cluster make predominantly own-payoff maximizing choices. We therefore label them as "predominantly selfish".

If our preference interpretation of the behavioral types is correct and stable across budget bundles, the different types should display characteristic behavioral patterns

²⁵Recall that subjects had to make a choice on twelve different budget lines. For each budget line, subjects could choose among seven different allocations. A choice is classified as own-payoff minimizing (own-payoff maximizing) if it belongs to the two choices that give the subject the lowest (highest) payoff. It is classified as payoff-equalizing if it implements perfect equality or one of its nearest neighbouring allocations.

in other situations (out of sample). For example, the inequality averse type should also display a preference for equality in new decision situations. Likewise, the selfish type should also predominantly maximize its own payoff in these alternative budget lines. In Fehr ^(r) al. (2022, 2023), we show that this is indeed the case.

Figure B.1: Distribution of choices for positively and negatively sloped budget lines in each cluster




C Additional Tables and Figures

C.1 Prior beliefs about income inequality

Several results hinge on conditioning on whether the respondent's income is below or above the median. It is therefore important to assess whether income predicts prior beliefs about the income share of the top 1%. Importantly, it does not. On average, respondents with an income below the median believe that the top 1% receives 54.4% (SD: 25.2 pp), while those with an income above the median believe that the top 1% receives 53.6% (SD: 26.0 pp), respectively. A Kruskal–Wallis equality-of-populations test cannot reject the null hypothesis of equal populations ($p = 0.665$). Prior beliefs do also *not* substantially differ when we further disaggregate the data by preference type and income category, as documented in Figure C.1. While some small differences exist, a Kruskal–Wallis equality-of-populations test cannot reject the null hypothesis of equality of populations ($p = 0.379$).

Figure C.1: Prior beliefs about the income share of the top 1%




Note: The figure shows the average prior beliefs about the income share of the top 1% of income-earners in Switzerland by income category and preference type (with standard errors).

C.2 Posterior beliefs about income inequality

Figure C.2 shows that the vast majority of the treated subjects (77.3%) correctly updated their beliefs, i.e., that they hold correct posterior beliefs.

Figure C.2: Posterior beliefs about the income share of the top 1% among subjects in the treatment group



Note: Distribution of posterior beliefs about the income share of the top 1% of income-earners in Switzerland among subjects in the treatment group. The vertical red line indicates the actual share of 12%.

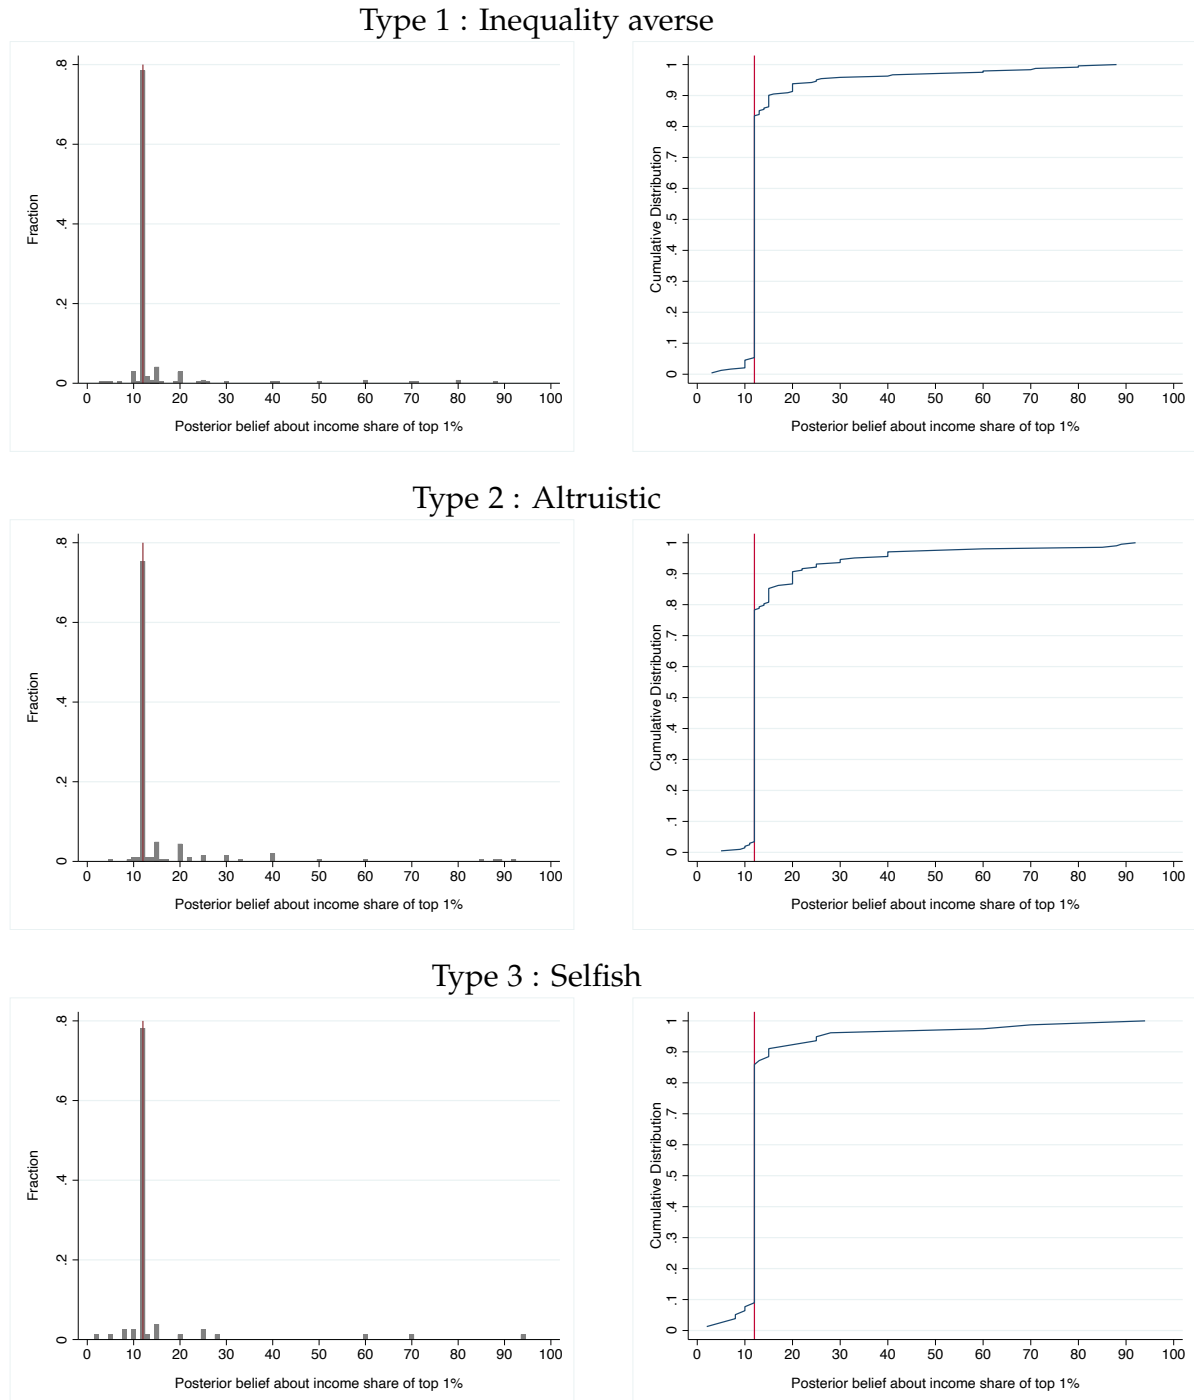
Importantly, individuals from all three preference types correctly update their beliefs, as shown in Section 5.3 of the main paper. Moreover, a Kruskal–Wallis test cannot reject the null hypothesis that individuals with different preference types update beliefs to a similar extent ($p = 0.686$). In addition, none of the Kolmogorov–Smirnov tests can reject the null hypothesis of equality of distributions when making pairwise comparisons of distributions in posterior beliefs.²⁶

Moreover, we also find conclusive evidence that respondents across preference types hold similar posterior beliefs about the income share of the top 1% (see Figure C.3). A Kruskal–Wallis test fails to reject the null hypothesis that individuals with different preference types hold similar posterior beliefs at conventional significant levels ($p = 0.088$). In addition, none of the Kolmogorov–Smirnov tests performed between preference types can reject the null hypothesis of equality of distributions of

²⁶Pairwise comparisons in belief updating (Kolmogorov–Smirnov tests): inequality averse vs. selfish ($p = 0.706$), altruistic vs. selfish ($p = 0.238$), inequality averse vs. altruistic ($p = 0.392$).

posterior beliefs.²⁷

Figure C.3: Posterior beliefs about the income share of the top 1%

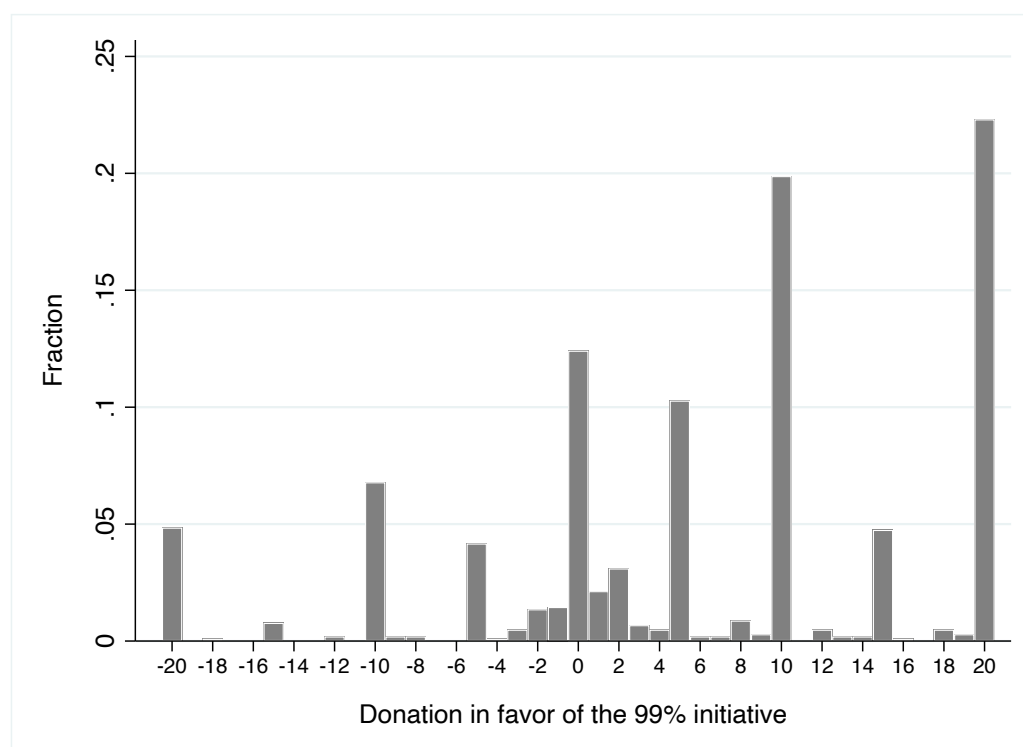


Note: Distribution of posterior beliefs about the income share of the top 1% of income-earners in Switzerland among subjects in the treatment group by preference type. The vertical red line indicates the actual share of 12%.

²⁷Pairwise comparisons of distribution of posterior beliefs (Kolmogorov-Smirnov tests): inequality averse vs. selfish ($p = 1.000$), altruistic vs. selfish ($p = 0.876$), inequality averse vs. altruistic ($p = 0.850$).

C.3 Donations

Figure C.4: Donations in favor of the 99% initiative



Note: Distribution of donation amount towards an organization in favor of the 99% initiative, with donations towards an organization that opposes the 99% initiative coded as negative values, i.e., the values can range from CHF -20 to CHF +20.

C.4 Full regression table

Table C.1: Determinants of donations in favor of the 99% initiative - full table

	Full sample				Below median income		Above median income	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	-1.106 (0.684)	-0.999 (0.668)	1.983 (1.881)	2.042 (1.825)	0.436 (2.694)	1.727 (2.815)	2.643 (2.860)	2.553 (2.718)
IA			6.871*** (1.519)	5.341*** (1.514)	4.721** (2.123)	4.501* (2.454)	8.232*** (2.288)	6.407*** (2.130)
Altruistic			5.001*** (1.554)	4.350*** (1.536)	3.777* (2.156)	3.730 (2.523)	5.133** (2.337)	5.434** (2.194)
Treat x IA			-4.298** (2.145)	-3.909* (2.091)	-1.001 (2.981)	-2.315 (3.077)	-7.252** (3.381)	-6.542** (3.247)
Treat x Altruistic			-2.951 (2.129)	-3.258 (2.066)	-0.716 (2.932)	-2.289 (3.069)	-4.252 (3.356)	-4.195 (3.266)
Belief of income share of top 1%		-0.003 (0.014)		-0.005 (0.014)		-0.019 (0.018)		0.010 (0.023)
Belief of own control over success (z)		-1.280*** (0.332)		-1.271*** (0.330)		-0.770* (0.439)		-1.764*** (0.575)
Belief of intergenerational income mobility (z)		-0.727** (0.353)		-0.622* (0.354)		-0.894* (0.469)		-0.606 (0.597)
Perceived degree of poverty in CH (z)		0.670* (0.349)		0.657* (0.351)		0.628 (0.454)		0.604 (0.645)
Believed gini of income distribution		3.855** (1.676)		4.137** (1.663)		5.382** (2.228)		4.242 (2.631)
Believed gini of wealth distribution		1.592 (1.922)		1.404 (1.916)		2.326 (2.443)		-1.743 (2.961)
Beliefs about future upwards mobility (z)		-0.412 (0.394)		-0.327 (0.399)		0.005 (0.483)		-1.365 (0.854)
Perceived past upwards mobility (z)		0.550 (0.398)		0.534 (0.396)		0.008 (0.468)		2.307*** (0.783)
Distrust in politicians (z)		0.034 (0.372)		-0.030 (0.364)		-0.385 (0.473)		0.613 (0.622)
Male		-0.160 (0.796)		0.057 (0.797)		-0.541 (1.029)		0.592 (1.534)
Age		0.140 (0.184)		0.147 (0.184)		0.025 (0.233)		0.576 (0.374)
Age squared		-0.001 (0.002)		-0.001 (0.002)		0.000 (0.003)		-0.005 (0.004)
French-speaking		0.478 (0.796)		0.492 (0.795)		0.488 (0.988)		1.470 (1.473)
Married		0.288 (0.750)		0.281 (0.746)		0.953 (1.033)		-1.033 (1.230)
Education: Vocational training		-0.949 (1.975)		-1.249 (2.128)		0.519 (2.600)		-4.720 (4.348)
Education: High school		-0.592 (2.114)		-0.970 (2.263)		0.804 (2.717)		-5.077 (4.804)
Education: University		-1.043 (2.018)		-1.359 (2.172)		0.608 (2.688)		-6.830 (4.313)
Education: Other		-1.187 (2.179)		-1.532 (2.321)		1.283 (2.858)		-7.722* (4.552)
Occupation: Part-time worker		1.188 (0.886)		1.327 (0.883)		-0.366 (1.249)		3.453** (1.555)
Occupation: Student		0.559 (1.375)		0.312 (1.392)		-1.883 (1.577)		-8.590** (3.834)
Occupation: Pensioner		2.397 (1.725)		2.678 (1.758)		1.155 (2.309)		4.439 (2.999)
Occupation: Unemployed		2.218 (1.793)		1.933 (1.798)		0.480 (2.029)		-1.363 (7.653)
Occupation: Other		-1.000 (1.593)		-0.738 (1.558)		-3.231* (1.873)		-0.115 (3.379)
History of being unemployed in the past		0.016 (0.730)		-0.058 (0.728)		-1.062 (0.922)		1.253 (1.323)
Risk aversion (z)		-0.364 (0.364)		-0.193 (0.364)		0.158 (0.463)		-0.201 (0.645)
Positive reciprocity (return favor)		1.058*** (0.368)		0.987*** (0.366)		0.831 (0.522)		0.965 (0.597)
Positive reciprocity (reciprocate help)		-0.235 (0.349)		-0.191 (0.348)		0.040 (0.464)		-0.581 (0.565)
Negative reciprocity (revenge unfair treatment)		-0.175 (0.526)		-0.176 (0.530)		0.365 (0.668)		-0.596 (0.876)
Negative reciprocity (retaliate intentional malice)		-0.242 (0.533)		-0.169 (0.539)		-0.089 (0.695)		-0.262 (0.896)
Impatience (z)		-0.223 (0.349)		-0.260 (0.344)		0.300 (0.454)		-0.698 (0.614)
Trust in people (z)		0.922** (0.367)		0.855** (0.359)		0.487 (0.490)		1.493** (0.617)
Wealth invested in financial assets		-1.129 (0.741)		-1.240* (0.734)		-0.432 (0.897)		-1.846 (1.297)
Constant	6.644*** (0.471)	-0.708 (4.592)	1.568 (1.376)	-4.795 (4.875)	3.314* (1.964)	-1.754 (6.243)	0.721 (1.996)	-11.661 (9.310)
Wealth bracket dummies	No	Yes	No	Yes	No	Yes	No	Yes
p-value(Ho: Treat x IA ≥ 0)			0.023	0.031	0.369	0.226	0.016	0.022
R ²	0.003	0.132	0.029	0.146	0.020	0.136	0.038	0.232
Observations	1031	1030	1031	1030	558	557	422	422

Notes. OLS regression. The dependent variable is the donation amount towards an organization in favor of the 99% initiative for the full sample (columns 1-4), for subjects with an income below the median (columns 5-6), and for subjects with an income above the median (columns 7-8). Subjects who did not disclose their income are not included in columns 5-8. Wealth bracket dummies include dummy variables for each wealth level. Robust standard errors are reported in parentheses. p-value(Ho: Treat x IA ≥ 0) indicates the p-value associated with a one-sided test evaluating the null hypothesis that the information intervention has a non-negative effect on the donations of the inequality averse subjects. Levels of significance: * $p < .1$, ** $p < .05$, *** $p < .01$.

C.5 Robustness analysis: attention checks

In this Appendix, we show that our main regression results are robust to excluding participants who did not successfully pass our attention checks. The second wave of the experiment (where the information intervention took place) included two attention checks. Table C.2 shows the regression results for participants who pass at least one attention check, and replicates the main results discussed in the paper. In particular, the interaction between the treatment and the inequality aversion dummy is significant for above-median income earners ($p = 0.014$ for $H_0 : \beta_4 \geq 0$). Table C.3 focuses only on subjects who successfully pass both attention checks and delivers a qualitatively similar message, although we are unfortunately underpowered to detect a significant interaction between the information intervention and the inequality aversion dummy for the respondents with an above-median income ($p = 0.068$ for $H_0 : \beta_4 \geq 0$).

Table C.2: Pass one of the two attention checks

	Full sample				Below median income		Above median income	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	-1.599** (0.705)	-1.391** (0.691)	1.553 (1.962)	2.039 (1.900)	-0.941 (2.875)	0.859 (3.027)	3.510 (2.898)	3.389 (2.758)
IA			7.280*** (1.602)	6.039*** (1.592)	5.105** (2.235)	5.486** (2.559)	8.647*** (2.435)	7.197*** (2.283)
Altruistic			5.421*** (1.633)	5.275*** (1.604)	4.005* (2.268)	4.372* (2.594)	6.093** (2.452)	6.649*** (2.321)
Treat x IA			-4.366* (2.240)	-4.487** (2.181)	-0.211 (3.165)	-2.023 (3.301)	-7.934** (3.477)	-7.331** (3.325)
Treat x Altruistic			-2.844 (2.205)	-3.460 (2.136)	0.405 (3.101)	-1.097 (3.269)	-5.579 (3.416)	-5.586* (3.305)
Constant	7.210*** (0.478)	-2.366 (4.479)	1.725 (1.465)	-6.832 (4.777)	3.484* (2.090)	-4.386 (6.131)	0.583 (2.117)	-12.865 (9.180)
Beliefs	No	Yes	No	Yes	No	Yes	No	Yes
Socio-demographics	No	Yes	No	Yes	No	Yes	No	Yes
Education	No	Yes	No	Yes	No	Yes	No	Yes
Occupation	No	Yes	No	Yes	No	Yes	No	Yes
Preference measures	No	Yes	No	Yes	No	Yes	No	Yes
Wealth bracket dummies & financial assets	No	Yes	No	Yes	No	Yes	No	Yes
p-value(Ho: Treat x IA \geq 0)			0.026	0.020	0.473	0.270	0.012	0.014
R ²	0.005	0.132	0.035	0.150	0.030	0.142	0.041	0.254
Observations	931	930	931	930	506	505	379	379

Note: OLS regression. The dependent variable is the donation amount towards an organization in favor of the 99% initiative for the full sample (columns 1-4), for subjects with an income below the median (columns 5-6), and for subjects with an income above the median (columns 7-8). Subjects who did not disclose their income are not included in columns 5-8. Beliefs include subjects' prior beliefs about the income share of the top 1%, their prior beliefs about income and wealth distributions in Switzerland, as well as their beliefs about the determinants of success, financial mobility, poverty in Switzerland, and distrust in politicians. Other socio-demographics include age, age squared, a dummy variable indicating whether the respondent is male, a dummy variable indicating whether the respondent's native language is French, and a dummy indicating whether the respondent is married. Education includes dummies indicating a respondent's highest educational achievement (compulsory school, vocational training, high school, university, or other). Occupation includes dummies indicating a respondent's occupation status (currently has a full-time job, a part-time job, is a student, is a pensioner, is unemployed, or other), a dummy variable indicating whether the respondent has experienced unemployment in the past. Preference measures from the global preference survey (Falk et al., 2022) include measures of risk aversion, (positive and negative) reciprocity, impatience, and general trust in strangers. Wealth bracket dummies and financial assets include dummy variables for each wealth level, and a dummy variable indicating whether the respondent has invested part of her wealth in financial assets. For a detailed explanation of the covariates, see Appendix A.1. Robust standard errors are reported in parentheses. p-value(Ho: Treat \times IA \geq 0) indicates the p-value associated with a one-sided test evaluating the null hypothesis that the information intervention has a non-negative effect on the donations of the inequality averse subjects. Levels of significance: * $p < .1$, ** $p < .05$, *** $p < .01$.

Table C.3: Pass both attention checks

	Full sample				Below median income		Above median income	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Treatment	-1.781** (0.749)	-1.717** (0.735)	1.435 (2.201)	1.897 (2.116)	-0.731 (3.336)	1.736 (3.485)	2.681 (3.103)	2.326 (2.849)
IA			7.128*** (1.820)	6.040*** (1.779)	4.621* (2.655)	5.289* (2.994)	8.185*** (2.702)	6.377** (2.513)
Altruistic			5.927*** (1.837)	5.627*** (1.774)	3.952 (2.677)	3.981 (2.986)	6.967*** (2.655)	7.264*** (2.576)
Treat x IA			-4.305* (2.485)	-4.645* (2.380)	-1.109 (3.622)	-3.722 (3.731)	-5.900 (3.741)	-5.258 (3.525)
Treat x Altruistic			-3.119 (2.437)	-3.702 (2.356)	-0.401 (3.544)	-2.397 (3.724)	-4.898 (3.644)	-4.212 (3.500)
Constant	7.302*** (0.514)	-5.215 (4.827)	1.667 (1.678)	-10.145* (5.183)	4.160* (2.524)	-7.497 (6.779)	0.033 (2.326)	-14.330 (9.413)
Beliefs	No	Yes	No	Yes	No	Yes	No	Yes
Socio-demographics	No	Yes	No	Yes	No	Yes	No	Yes
Education	No	Yes	No	Yes	No	Yes	No	Yes
Occupation	No	Yes	No	Yes	No	Yes	No	Yes
Preference measures	No	Yes	No	Yes	No	Yes	No	Yes
Wealth bracket dummies & financial assets	No	Yes	No	Yes	No	Yes	No	Yes
p-value(Ho: Treat x IA \geq 0)			0.042	0.026	0.380	0.159	0.058	0.068
R ²	0.007	0.137	0.035	0.155	0.024	0.147	0.038	0.262
Observations	817	816	817	816	449	448	329	329

Note: OLS regression. The dependent variable is the donation amount towards an organization in favor of the 99% initiative for the full sample (columns 1-4), for subjects with an income below the median (columns 5-6), and for subjects with an income above the median (columns 7-8). Subjects who did not disclose their income are not included in columns 5-8. Beliefs include subjects' prior beliefs about the income share of the top 1%, their prior beliefs about income and wealth distributions in Switzerland, as well as their beliefs about the determinants of success, financial mobility, poverty in Switzerland, and distrust in politicians. Other socio-demographics include age, age squared, a dummy variable indicating whether the respondent is male, a dummy variable indicating whether the respondent's native language is French, and a dummy indicating whether the respondent is married. Education includes dummies indicating a respondent's highest educational achievement (compulsory school, vocational training, high school, university, or other). Occupation includes dummies indicating a respondent's occupation status (currently has a full-time job, a part-time job, is a student, is a pensioner, is unemployed, or other), a dummy variable indicating whether the respondent has experienced unemployment in the past. Preference measures from the global preference survey (Falk et al., 2022) include measures of risk aversion, (positive and negative) reciprocity, impatience, and general trust in strangers. Wealth bracket dummies and financial assets include dummy variables for each wealth level, and a dummy variable indicating whether the respondent has invested part of her wealth in financial assets. For a detailed explanation of the covariates, see Appendix A.1. Robust standard errors are reported in parentheses. p-value(Ho: Treat \times IA \geq 0) indicates the p-value associated with a one-sided test evaluating the null hypothesis that the information intervention has a non-negative effect on the donations of the inequality averse subjects. Levels of significance: * $p < .1$, ** $p < .05$, *** $p < .01$.