NETWORK LISER

Discussion Paper Series

IZA DP No. 19037

October 2026

The Role of Social Preferences for Life Goals and Career Choices

Julien Senn ⓡ
University of Paris 1 Pantheon-Sorbonne

Thomas F. Epper ⓡ
CNRS, IESEG School of Management
and iRisk Center on Risk and Uncertainty

Ernst Fehr
University of Zurich
and IZA@LISER

The IZA Discussion Paper Series (ISSN: 2365-9793) ("Series") is the primary platform for disseminating research produced within the framework of the IZA@LISER Network, an unincorporated international network of labour economists coordinated by the Luxembourg Institute of Socio-Economic Research (LISER). The Series is operated by LISER, a Luxembourg public establishment (établissement public) registered with the Luxembourg Business Registers under number J57, with its registered office at 11, Porte des Sciences, 4366 Esch-sur-Alzette, Grand Duchy of Luxembourg.

Any opinions expressed in this Series are solely those of the author(s). LISER accepts no responsibility or liability for the content of the contributions published herein. LISER adheres to the European Code of Conduct for Research Integrity. Contributions published in this Series present preliminary work intended to foster academic debate. They may be revised, are not definitive, and should be cited accordingly. Copyright remains with the author(s) unless otherwise indicated.

IZA DP No. 19037 October 2026

The Role of Social Preferences for Life Goals and Career Choices*

Abstract

Workers care not only about pay and career opportunities but also about meaningful work and socially responsible employers. This heterogeneity matters for worker-firm matching, recruitment, and compensation, yet its preference-based origins remain poorly understood. In a two-wave study of a sample broadly representative of the Swiss adult population (N=1,031), we study whether social preferences provide a common account of life goals, job preferences, and occupational sorting. We structurally estimate two facets of social preferences: altruism in the advantageous and in the disadvantageous domain. We show that these two dimensions are negatively correlated, and that failing to account for both can substantially bias — in our applications, underestimate — the effects of social preferences. Social preferences predict respondents' life goals, their preferred job attributes, and their willingness to work in the tobacco industry. In particular, higher altruism is associated with lower emphasis on wealth, financial security, career status, pay, and advancement, but greater emphasis on corporate social responsibility and meaningful work. It also predicts greater willingness to avoid working in the tobacco industry. Lastly, altruism when ahead is positively associated with employment in education, health, and social services. Thus, this paper highlights the importance of social preferences for explaining worker motivation and labor-market sorting.

JEL classification

D01, D91, J24

Keywords

social preferences, meaning-at-work, career choices, sorting

Corresponding author

Ernst Fehr
ernst.fehr@econ.uzh.ch

* We thank Simon Gaechter, Rafael Lalive, Michel Maréchal, Christian Zehnder, Florian Zimmermann for helpful comments and suggestions. Thomas Epper gratefully acknowledges the financial support received from the Métropole Européenne de Lille (MEL). Ethics approval was obtained from the Human Subjects Committee of the Department of Economics of the University of Zurich (OEC IRB #2020-093 and #2021-032).

1 Introduction

People's aspirations and life goals may influence important dimensions of their labor supply. For some people, money, wealth, and making a career appear to be very important for overall life satisfaction. For others, non-pecuniary aspects may play a more important role. For example, previous work has shown that some people care a lot about doing work that is meaningful (see, e.g., Carpenter and Gong, 2016; Cassar, 2019; Non et al., 2022), some value working for a company that is socially responsible (see, e.g., Schneider et al., 2024; Burbano, 2016; Non et al., 2022), and others care about flexibility (Mas and Pallais, 2017, 2020). While previous work has shown these dimensions may affect selection at work (Burbano et al., 2024; Schneider et al., 2024; Cassar and Meier, 2018), relatively little is known about the origins of such preferences.

In this paper, we examine whether social preferences provide a common preference-based account of heterogeneity in life goals, demand for job-attributes, and career choices. We are not the first to hypothesize a link between social preferences and labor market outcomes. However, existing evidence mainly relies on non-incentivized measures of social preferences (Kosse and Tincani, 2020; Dur and Zoutenbier, 2014; Dohmen et al., 2009) or on relatively coarse proxies for social preferences such as a single dictator game (Kolstad and Lindkvist, 2013; Lagarde and Blaauw, 2014; Carpenter and Myers, 2010) where the donor has a distributional advantage and the sacrifice of $1 generates a benefit of $1 to the recipient, a trust game (Gill et al., 2023; Barr and Serneels, 2009), or donations to charities (Buurman et al., 2012).

An important limitation of proxying social preferences using a single, standard dictator game or charitable donations is that they only allow us to identify one form of prosociality in a rather coarse way: altruism in the advantageous domain (i.e., when ahead of the other) at a single price of giving. Yet, a growing body of empirical evidence suggests that individuals' social preferences may be a lot more diverse and require understanding how individuals respond to the cost of other-regarding actions (see, e.g. Fehr ⓡ al., 2024, 2023; Fisman et al., 2017; Kerschbamer and Müller, 2020).[1]

[1]For example, Fehr ⓡ al. (2024, 2023) document the existence of 3 types: (i) a majority of individuals who are willing to behave altruistically when ahead, but reduce others' payoff when behind, (ii)

In this paper, we study the nature and the role of a broader and more finely measured set of social preferences for people's preferences over life goals, job-attributes, and career choices. We formalize a subject's utility in a two-person case as a linear combination of the decision maker's payoff and the payoff of an anonymous recipient. Crucially, we allow the weight on the other participant's payoff to change depending on whether the decision maker is better or worse off than the recipient. This formulation, which nests inequality aversion (Fehr and Schmidt, 1999) and Charness-Rabin-type distributional preferences (Charness and Rabin, 2002), has the advantage that it lets us describe behavior directly in terms of an individual's degree of altruism in the advantageous and the disadvantageous domain. To identify these preference weights more precisely, we elicit choices under several different costs of redistribution and estimate them structurally at the individual level. This allows us to capture fine-grained differences in the strength of altruism, including differences between individuals who make similar choices at a one-to-one price of giving but respond differently when giving becomes more or less costly. With this framework in mind, we ask the following questions:

1. What is the empirical link between altruism in the advantageous domain and altruism in the disadvantageous domain?

2. What are the implications of accounting *only* for altruism in the advantageous domain, and neglecting altruism in the disadvantageous domain?

3. Can social preferences provide a common preference-based account of heterogeneity in life goals, demand for job-attributes, and career choices?

We study these questions using a two-wave study with a sample that is broadly representative of the Swiss adult population. In the first wave, we elicit participants' social preferences using an incentivized money allocation task in which participants have to make distributional choices between themselves and another anonymous participant. We ask subjects to make a decision in a large variety of choice situations in

a sizable share of individuals who are not only willing to behave altruistically when ahead, but also when behind, and (iii) a sizable minority of selfish individuals who predominantly care about maximizing their own payoff. Fisman et al. (2017) show that the price of giving is a core determinant of other-regarding behavior and that individuals differ according to their willingness to give at various prices.

which we systematically vary the cost of redistribution, thereby allowing us to identify a much wider range of preference types than what would be possible using a single dictator game.

In the second wave, conducted about six months later, we measure various aspects of participants' preferences that might influence their supply of labor. In one battery of questions, we ask about their *life goals*, which we define as what respondents value in life, i.e., how different aspects of their life contribute to their overall life satisfaction (Zukin and Szeltner, 2012). Our measures cover pecuniary aspects, such as being financially secure, being wealthy, or having a prestigious career, as well as non-pecuniary aspects such as being married, or having children. In another battery of questions, we measure respondents' preferences over job-attributes. The goal of these questions is to understand how individuals make trade-offs between the various aspects of a job. We are particularly interested in the weight people put on i) doing work that is meaningful, and ii) working for a company that values social and environmental responsibility—two dimensions that have been shown to play an important role for the labor supply of individuals (see, e.g., Cassar and Meier, 2018). In addition, we also measure the weight that people put on various pecuniary aspects of a job (e.g., high income and chances for advancement), as well as other job characteristics such as flexibility (Mas and Pallais, 2017, 2020) and short working hours. Together, this study wave helps reveal what people value *in life* and *at work*. We complement it with a simple discrete choice experiment in which we measure respondents' willingness to *give up* money to avoid working in the tobacco industry—an industry typically perceived as immoral (Schneider et al., 2024). In addition, our surveys also include a large number of questions on participants' sociodemographics such as age, gender, education, income, and related characteristics that will serve as controls.

This staggered design, which decouples the elicitation of respondents' social preferences from the possible determinants of their labor supply, allows us to explore the extent to which social preferences can offer a common preference-based account for individuals' heterogeneity in life goals, job-attribute preferences, and career-related outcomes. We complement this survey-based data with third-party data on respondents' industry of occupation. We use this data to get insights about participants'

sorting in the labor market.

We start our analysis by exploring the nature of social preferences in our sample. To that end, we estimate a structural model of altruism. We find that subjects display a substantial average willingness to give when ahead, i.e., they behave in an altruistic way when they are in an advantageous situation. At the same time, however, they also display on average a *negative* willingness to give (i.e., a willingness to reduce the other person's income) when behind—that is, they do *not* behave in an altruistic way when they are in a disadvantageous position. There is, however, substantial heterogeneity among the subjects: A considerable share of participants does not display a willingness to reduce the partner's payoff when behind (i.e., is not behindness averse), but instead display altruism in the disadvantageous domain.

In addition, we find a significant negative Spearman rank correlation of $\rho = -0.507$ between altruism in the advantageous and disadvantageous domains: individuals who are more altruistic when ahead tend to place a less positive (or more negative) weight on the other person's payoff when behind. To clarify the meaning of this correlation, consider individuals who are altruistic in the advantageous domain and distinguish between two cases: (i) Those who are also altruistic in the disadvantageous domain. For them, the negative correlation means that greater altruism when ahead is associated with weaker altruism when behind. (ii) Those who are inequality averse, i.e., they are altruistic when ahead but behindness averse in the disadvantageous domain: they place a negative weight on the other person's payoff when behind and are therefore willing to sacrifice own income to reduce the other person's income. For them, the same negative correlation means that greater altruism when ahead is associated with stronger behindness aversion when behind. Thus, the negative correlation means, e.g., that individuals who display stronger altruistic behavior in the standard dictator game may simultaneously also be more behindness averse.

We then explore the link between social preferences, life goals, preferences over job-attributes, and other career-related choices. A possibly important implication of this (negative) correlation between altruism when ahead and when behind is that omitting to account for both dimensions of altruism can substantially bias the effects of social preferences. In fact, we show that including only altruism when ahead—which is what, to the best of our knowledge, most existing studies typically do—

yields effects of social preferences that are biased towards zero, i.e., the effects of social preferences tend to be underestimated when we do not account for both dimensions of altruism. We show that this result—which can be analytically derived from the omitted variable bias formula—depends on i) the sign of the effect of these two dimensions of altruism for the outcome of interest, and ii) the sign of the correlation between these two dimensions of altruism.

In addition to the negative correlation between altruism in the advantageous and the disadvantageous domain, we document several other new results. First, we find that altruism is a sizable predictor of the importance that subjects attribute to career and financial aspects for their overall life satisfaction. For example, people with greater altruism in both the advantageous and the disadvantageous domain deem having a prestigious career, and financial security, as *less* important for their overall life satisfaction. For example, a one-standard-deviation increase in altruism in the advantageous (disadvantageous) domain leads to a 0.089 (0.106) standard deviation decrease in subjects' assessments of the importance of having a prestigious career while risk aversion and patience play no role. The coefficients on both dimensions of altruism are sizable: their magnitude is about 50% to 60% of the effect size of the gender dummy. Moreover, a one-standard-deviation increase in altruism in the advantageous (disadvantageous) domain is associated with a 0.118 (0.090) standard deviation decrease in the importance of financial security. This is an even higher (about the same) effect size as the 0.088 standard deviation decrease associated with an increase in the willingness to take risks. Finally, we find that a one-standard-deviation increase in altruism in the advantageous domain yields a 0.155 decrease in subjects' assessment of the importance of being wealthy for their overall life satisfaction. Overall, people who are more altruistic rate the pecuniary and career-related items as being *less* important for their overall life satisfaction than less altruistic respondents. Interestingly, altruism is a poor predictor of non-job-related items like having a partner or having children for overall life satisfaction.

If people with stronger altruism deem career and the pecuniary aspects of their life as less important for their overall life satisfaction, do they also prioritize jobs with different attributes and are they more likely to select into different industries? We find that people with a higher altruism rank working for a company that values

social and environmental responsibility systematically higher. This effect is sizable: the effect of altruism in the advantageous domain on the rank assigned to working for a socially responsible company is comparable to the difference between females and males. In addition, we also find that people who are more altruistic when ahead attach significantly more importance to doing work that is meaningful for society. However, more altruistic respondents attach significantly less importance to the more pecuniary aspects of a job like pay, financial security and chances for advancement—which is consistent with the finding that those with higher altruism deem financial considerations less important for their overall life satisfaction.

Finally, we examine whether social preferences are also related to respondents' willingness to make monetary tradeoffs in a stylized labor-market choice experiment and to their actual occupational sorting. In the choice experiment, respondents with stronger social preferences are willing to forgo more salary to avoid working for the tobacco industry. Using independently collected occupation data, we also find that altruism in the advantageous domain is positively associated with employment in social jobs such as education, health, and social services. This latter relationship is quantitatively more modest, but it provides complementary evidence that experimentally elicited social preferences are related not only to stated goals and job-attribute rankings, but also to observed labor-market outcomes.

Taken together, our findings suggest that social preferences provide a common preference-based account of systematic heterogeneity in what people value in life and at work. More altruistic individuals attach less importance to career status, wealth, financial security, and pecuniary job-attributes, while placing greater weight on corporate social responsibility and work that is meaningful for society. At the same time, our results underscore an important methodological point: Approaches that do not separately measure and jointly account for altruism in the advantageous and the disadvantageous domain may miss important structure in social preferences and generate biased estimates of their relationship with economically relevant outcomes.

2 Related Literature

Our paper relates to several strands of the literature.

Social preferences and sorting at the workplace First, our paper connects to the literature linking social preferences to labor market outcomes and sorting. Early contributions have explored the link between altruism and various labor market outcomes such as, e.g., the labor supply of volunteer firefighters (Carpenter and Myers, 2010), the selection of health worker students into the public sector (Kolstad and Lindkvist, 2013; Dur and Zoutenbier, 2014), and geographical sorting (Lagarde and Blaauw, 2014). Other papers have focused on the role of reciprocity (see, e.g., Dohmen et al., 2009; Barr and Serneels, 2009). More recently, several contributions have explored the role of trust(worthiness) and honesty for selection into the public sector (Hanna and Wang, 2017; Barfort et al., 2019), the police force (Friebel et al., 2019), the financial industry (Gill et al., 2023), and immoral industries (Schneider et al., 2024).[2] However, as we noted above, much of this literature relies on non-incentivized measures of social preferences (Kosse and Tincani, 2020; Dur and Zoutenbier, 2014; Dohmen et al., 2009) or on relatively coarse proxies for social preferences such as a single, standard dictator game (Kolstad and Lindkvist, 2013; Lagarde and Blaauw, 2014; Carpenter and Myers, 2010), trust games (Gill et al., 2023; Barr and Serneels, 2009) or donations to charities (Buurman et al., 2012).

Perhaps most closely connected to our approach are the recent papers analyzing the distributional preferences of elite students (Fisman et al., 2015), medical students (Li et al., 2017) and physicians (Li et al., 2022). While these papers also approach the elicitation of social preferences using incentivized choices, their design is restricted to choice situations in which subjects can give up some of their endowment in order to *increase* the payoff of the recipient, but does not include choice situations where the subjects can sacrifice money in order to *decrease* the payoff of the recipient. As a consequence, their design does not identify a wide range of preferences such as envy, spite, or aversion to disadvantageous inequality.

Against this background, an important contribution of our paper is to complement existing approaches to measuring social preferences with an experimental design that identifies a broader range of distributional preferences. This richer measurement framework enables us to document the negative correlation between altruism in

[2]See also Kosse and Tincani (2020), who explore the relation between prosociality and labor market success in a large cross-country survey.

the advantageous and disadvantageous domains and to account for both dimensions jointly in our analysis. As we discuss, the lack of a simultaneous control for both of these facets of altruism yields strongly biased estimates, suggesting that the effects of social preferences might have been underestimated in previous studies. In addition, most existing work focuses on the link between social preferences and (self-reported) selection in the labor-market. We expand the analysis and show that social preferences also offer an account for respondents' life goals and their preferences over job-attributes.

Meaning-at-work and job mission Second, our paper connects to the growing literature on preferences over job-attributes, non-monetary incentives and the implications of work as a source of meaning (Cassar and Meier, 2018; Henderson and Steen, 2015; Gartenberg et al., 2019; Chater and Loewenstein, 2016; Karlsson et al., 2004; Rosso et al., 2010; Wrzesniewski et al., 2003). For example, recent contributions have explored the role of meaning and job mission for worker motivation (Cassar, 2019; Carpenter and Gong, 2016), job attractiveness (Burbano, 2016), sorting (Non et al., 2022), and optimal contracting (Cassar and Armouti-Hansen, 2020), among others. An important open question in this literature is to better understand what explains cross-individual heterogeneity in such preferences over meaning. While the latest contributions have highlighted the role of macroeconomic conditions during impressionable years (Cotofan et al., 2023) and gender (Burbano et al., 2024) as important predictors, our paper offers a preference-based explanation that can jointly account for heterogeneity in life goals, preferences over job-attributes, and sorting into the labor market.

External validity of laboratory measures of social preferences Last, our paper connects to the recent literature exploring the predictive ability of social preferences for real world behavior. Among others, experimentally elicited measures of social preferences have been shown to be significant predictors of a host of behaviors such as demand for redistribution (Henkel ⓡ al., 2025; Fehr ⓡ al., 2024; Fisman et al., 2015; Kerschbamer and Müller, 2020; Epper et al., 2024), health behaviors (Campos-Mercade et al., 2021; Schneider et al., 2023; Alfaro et al., 2022; Müller and Rau, 2021;

Galizzi and Wiesen, 2018), investment decisions (Riedl and Smeets, 2017; Bauer et al., 2021), honesty (Cohn and Maréchal, 2018; Dai et al., 2018), and decisions involving environmental externalities (Campos-Mercade et al., 2025). Our paper adds a rich set of labor market-related outcomes to this list, and also highlights the importance of accounting for both altruism in the advantageous and the disadvantageous domain.

3 Experimental design and data description

3.1 Overview

Our aim is to study the link between people's social preferences, their life goals, their preferences over job-attributes and other career-related preferences. In particular, we are interested in understanding whether social preferences predict people's preferences over job-attributes, such as higher pay, mission, or working for a company that values social and environmental responsibility.

This objective requires the following ingredients: (i) a rich measure of respondents' social preferences, (ii) *independent* measures of their life goals and their views on different aspects of their (preferred) professional career, (iii) a large array of individual controls. We approach this task with a two-wave survey experiment.

The main objective of the first wave is to measure subjects' social preferences. As we detail in the next subsection, we approach this task by relying on a large set of incentivized dictator games in which subjects have to allocate money between themselves and an anonymous recipient. In addition, we also use this opportunity to collect information on participants' socio-demographics. We also include a validated survey module aimed at measuring other economic preferences (Falk et al., 2016). In particular, we will use these measures to control for subjects' time and risk preferences.

In the second wave—which we conducted about six months later—we collected detailed information on respondents' life goals as well as their preferences over job-attributes. We also included a hypothetical choice experiment aimed at gauging respondents' (un)willingness to work for the tobacco industry—an industry typically perceived as immoral (see, e.g., Schneider et al., 2024).

We complement this survey with independently collected third-party data on respondents' occupation, which allows us to explore the link between respondents' social preferences and their selection into different types of occupations. In the following, we provide details on the main parts of both study waves, the third-party dataset, as well as the subject pool.

We summarize the key features of each wave below. For details, see Appendix A.

3.2 Wave 1: Eliciting social preferences

We conducted the first wave in December 2020. The aim of this wave is to uncover the distribution of other-regarding preferences of our respondents. We approach this task by eliciting their distributional preferences in a set of 20 incentivized money allocation tasks.[3] In each of these tasks, the participants decide how to allocate experimental currency units (ECUs) between themselves and an anonymous other participant of the study.

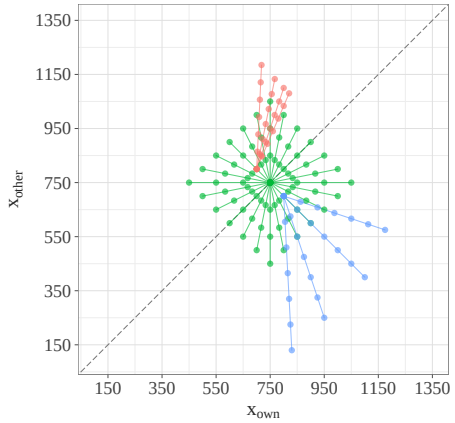
Figure 1a depicts the various allocation lines for which subjects had to make a decision. In some decision situations, the decision maker can give up some of her own payoff to *increase* the payoff of the other. In other decision situations, they can pay to *decrease* the payoff of the other. These different 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.

Figure 1b 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 an allocation 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 allocation 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.

[3]The money allocation task used for the elicitation of social preferences is based on Fehr ⓡ al. (2024, 2023).

Figure 1: Measuring other-regarding preferences with a money allocation task

(a) Allocation lines



(b) Decision screen



Note: Figure 1a depicts the 20 decision situations (allocation lines) for which subjects had to make a decision. For each allocation line (i.e., for each decision situation), subjects were asked to choose an interpersonal allocation of ECUs between themselves (x_{own}) and the other person (x_{other}). Each "budget line" comprised seven different discrete interpersonal allocations, represented by dots in Figure 1a. Figure 1b illustrates how a typical choice situation was presented to participants.

3.3 Wave 2: Life goals and career-related preferences

About six months later, we contacted the *same* subjects to participate in a new study. This study wave included a number of questions aimed at understanding what matters most for respondents' overall life satisfaction (i.e., their life goals) and their preferences over different job-attributes–among others. In particular, this wave included the following modules, which are at the heart of our analysis in this paper.

Life goals. One set of questions aims at understanding what respondents value in life, i.e., how different aspects of their life, ranging from having children to being financially secure, contribute to their overall life satisfaction. Specifically, we ask the following battery of questions (inspired by Zukin and Szeltner, 2012):

> *Different people have different goals in life. Please indicate how important each of the following is to your overall life satisfaction. [Having a partner / being married; Being wealthy; Having children; Having a prestigious career; Being financially secure; Being a leader in my community; Having a job where I can make an impact on causes or issues that are important to me]*

For each item, participants answer on a 1 to 5 scale where 1 means *"Not at all im-*

portant" and 5 means *"Very important"*. We use these questions to gauge what we refer to as people's *life goals and aspirations*—the aspects that our respondents believe contribute the most to their overall life satisfaction.

Preferences over job-attributes Another battery of questions aims at understanding what kind of job-attributes our respondents value. While some respondents might mainly care about pay and career advancement, others might find it more important to work for a company that is socially responsible (e.g., Burbano, 2016; Schneider et al., 2024; Non et al., 2022), to have a job that is meaningful (e.g., Cassar and Meier, 2018; Carpenter and Gong, 2016; Cassar, 2019), or to have flexibility in their work arrangements (Mas and Pallais, 2020, 2017), among others. To assess the extent to which respondents trade off different aspects of a job, we ask respondents to *rank* a variety of job-attributes, from the most important (ranked 1st) to their least preferred one (ranked last). This battery of questions is based on the General Social Survey.

> *Please look at the list below and tell us which one thing on this list you would most prefer in a job? Which comes next? Which is third-most important, etc.? [High income; Job Security; Working hours are short, lots of free time; Flexibility in work hours; Chances for advancement; Work is meaningful for (useful to) society; Working for a company that values social and environmental responsibility; A work environment that allows to form friendships with colleagues]*

(Un)willingness to work for the tobacco industry Previous work has shown that workers can be reluctant to work in industries that are perceived as immoral (see, e.g., Schneider et al., 2024). In order to gauge our respondents' (un)willingness to work for such industries, our study also included the following hypothetical choice experiment, inspired by Frank (1996):

Imagine that you are looking for a new job and that two companies offered you a job. Company A is a large international tobacco company. Company B is a medical research company which helps to develop new cures for diseases. Both companies will give you some control over what you do and how you do it, but in both companies your boss will tend to set rather strict guidelines. Company A offers you a monthly salary of CHF 5,900. You are still unsure about the monthly salary that will be offered to you by company B. For each of the possible monthly salaries that company B might offer you, please tell us whether you prefer to take the job offer from company A or the job offer from company B.

The range of possible salaries offered by the medical company spanned from CHF 5,000 to CHF 6,700 and we enforced a single switching point. This choice experiment allows us to assess whether and how much salary respondents are willing to give up to avoid working for the tobacco company.[4,5]

3.4 Occupation data and other survey measures

In addition, we also collected the following data.

Occupation data We obtained *third-party* data on respondents' industry of occupation. This data was collected independently of us by our survey provider. It contains information on respondents' main occupation at the time of our study, with 25 categories of occupations such as *"Transports"*, *"Construction"*, *"Education"*, or *"Health services"*.

Sociodemographics and further covariates In addition, we collected an extensive amount of information on participants' sociodemographics such as age, gender, education, income, etc. We also measured other facets of economic preferences, in particular time and risk preferences, using the experimentally validated preference

[4]For example, if a respondent prefers to work for the medical company at a salary of CHF 5,000 instead of working for the tobacco company at a salary of CHF 5,900, this means that she is willing to give up CHF 900 in income (about 15 percent) to avoid working for the tobacco company.

[5]Other papers have used related methods to estimate people's valuation for different job characteristics. See, e.g., Burbano et al. (2024), Non et al. (2022), and Wiswall and Zafar (2018)

survey module (Falk et al., 2016).[6]

3.5 Data collection, sample and experimental protocol

Ethics approval was obtained from the Human Subjects Committee of the Department of Economics of the University of Zurich (OEC IRB #2020-093 and #2021-032). We collaborated with the LINK Institute for data collection. The study was conducted online with a sample that is broadly representative of the German- and the French-speaking adult population of Switzerland.[7] 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.[8] A total of 1,383 subjects completed wave 1, while 1,031 subjects completed wave 2. In our analyses below, we focus on the 1,031 subjects who completed both waves. Descriptive statistics on participants' main socio-demographic characteristics can be found in Table S.1 in Appendix B. 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%, the median income is CHF 6,000, and the share of unemployed respondents is 3%.

Attrition. An important asset of our study design is that the elicitation of social preferences is decoupled from the survey block on life goals and career-related preferences. The advantage of such a staggered design is that it limits demand effects and concerns that subjects might want to behave consistently across tasks. For exam-

[6]Both study waves comprised a number of additional, unrelated, survey blocks. The first wave included survey items aimed at measuring respondents' beliefs about inequality, social mobility, poverty, determinants of success, and trust in politicians. Wave 2 included a number of questions aimed at understanding respondents' beliefs about inequality and their support for redistribution, as well as an information intervention, which we reported on in a separate paper (Henkel ⓡ al., 2025). We discuss these survey items in greater detail in Appendix A.1 and A.2. In addition, we provide evidence that excludes cross-study spillovers in Appendix B.3.

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

[8]Median 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).

ple, a subject who shares in the dictator game might want to appear consistent and indicate a preference for working in a company that values social and environmental responsibility if that question was asked directly after the dictator game. A staggered experimental design mitigates concerns about such cross-task spillovers. At the same time, however, a potential risk of multi-wave studies is selective attrition. In Table S.3 in Appendix B.2, we show that attrition between the two waves is orthogonal to our social preferences estimates, and to the bulk of participants' observable characteristics, i.e., we have very little reason to worry about selective attrition in our sample.

Limited attention. A possible problem with online studies is that some of the subjects might not pay attention to the questions that are asked to them, which might affect the quality of the results. To address this possible issue, we added two attention checks that are aimed at measuring whether participants read survey items carefully before answering (Berinsky et al., 2014), i.e., they allow us to measure respondents' attentiveness and to proxy data quality. Data quality in our sample is remarkably high, with 79.2% of the subjects correctly answering both attention checks, and only 9.7% failing to pass both checks.[9]

4 The empirical distribution of social preferences

4.1 Preference estimation

We elicit subjects' distributional preferences through a randomized sequence of allocation decisions. In each decision, subject i chooses one of seven interpersonal allocations located on a line segment in the own-other payoff space, where i denotes the subject and j denotes the other person.

Let (x_i, x_j) denote a feasible allocation, where x_i is the decision maker's payoff and x_j is the other person's payoff. The allocation lines in our design vary in both slope and position (Figure 1a), so subjects encounter choices in which they are ahead, behind, or can choose on either side of equality. This variation allows us to study

[9]Some online samples contain considerable proportions of respondents who do not correctly answer attention check questions even in shorter surveys, with failure 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.

giving behavior separately under advantageous and disadvantageous inequality.

We model subject i's utility over allocations (x_i, x_j) using the following piecewise-linear specification:

$$U_i(x_i, x_j) = \begin{cases} (1-\beta_i)x_i + \beta_i x_j & \text{if } x_i \geq x_j, \\ (1-\alpha_i)x_i + \alpha_i x_j & \text{if } x_i < x_j. \end{cases} \tag{1}$$

The parameter β_i captures subject i's marginal concern for the other participant's payoff when i is ahead, whereas α_i captures the corresponding marginal concern when i is behind. This specification is algebraically equivalent to the Fehr and Schmidt (1999) model, except that we reverse the sign of the disadvantageous-inequality parameter, α_i. Under this convention, negative values of α_i indicate aversion to disadvantageous inequality, while positive values indicate altruistic concern for the other person even when behind.

To account for stochastic choices, we embed the utility specification in a logit choice model. For each candidate parameter vector $\theta_i = (\alpha_i, \beta_i, \lambda_i)$, the model assigns a utility value to every feasible allocation in the choice set and maps these utilities into choice probabilities:

$$P_i\big(x_i, x_j \mid A, \theta_i\big) = \frac{\exp\big(U_i(x_i, x_j)/\lambda_i\big)}{\sum_{(y_i, y_j)\in A} \exp\big(U_i(y_i, y_j)/\lambda_i\big)}. \tag{2}$$

Here, A denotes the set of feasible allocations in a given decision problem. The noise scale $\lambda_i > 0$ governs how closely choices follow utility maximization: as λ_i approaches zero, choices become nearly deterministic; larger values of λ_i imply noisier choices.

Our design is informative because each allocation line implies cutoff values for the altruism parameters. For a given line, its slope and position determine the parameter values at which the utility-maximizing allocation changes. Lines that lie entirely in the disadvantageous domain, where $x_i < x_j$, identify restrictions on α_i because utility depends only on the α_i parameter along those lines. Lines that lie entirely in the advantageous domain, where $x_i \geq x_j$, identify restrictions on β_i. Lines that cross the equality line, where $x_i = x_j$, are informative about both α_i and β_i, since feasible allocations on the same line may fall on either side of equality. Thus, observed choices

locate each subject's parameters relative to a set of line-specific cutoffs. This cutoff structure is the main source of identification in our design.

We estimate the model separately for each subject using a Bayesian approach. We construct a grid over (α_i, β_i) from the design-implied cutoff structure and combine it with an exogenous log-spaced grid for λ_i. For each subject, we evaluate the likelihood of the observed choice sequence at every grid point and combine it with a uniform prior over the grid intervals.[10] This yields an individual-level posterior distribution over $(\alpha_i, \beta_i, \lambda_i)$.

In what follows, we use the individual posterior means of the altruism parameters as subject-level summaries. The corresponding posterior distributions indicate how precisely each subject's parameters are identified. Appendix D provides details on the estimation approach.

4.2 Structural estimates

We present the summary statistics for the individual posterior mean parameters in Table 1 below. The average level of altruism when behind (α) is -0.287 (with a standard deviation of 0.605), and the average level of altruism when ahead (β) is 0.874 (with a standard deviation of 0.933). This implies that, on average, subjects value the other's payoff positively when ahead, but negatively when behind.

Table 1: Summary statistics of structural estimates

	Mean	S.D.	Min	Max
Altruism when behind (α)	-0.287	0.605	-1.125	2.374
Altruism when ahead (β)	0.874	0.933	-1.375	2.375
Observations	1031			

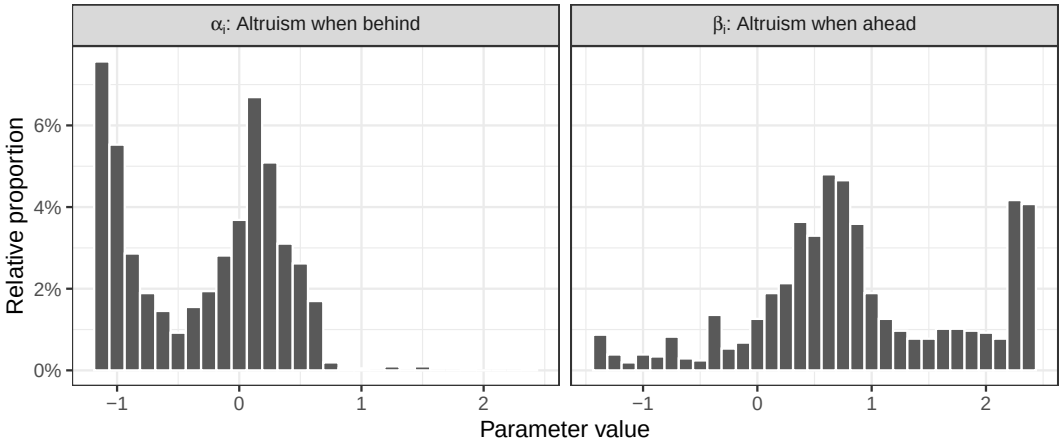
In Figure 2, we depict the distributions of the individual mean α_i and β_i. The figure indicates that there is wide individual-level heterogeneity in both dimensions of altruism. The left panel shows the distribution of individual posterior means of altruism in the disadvantageous domain (α_i). This individual posterior mean of α_i

[10] In Appendix D, we show that our estimation results are robust to specifying the prior as uniform over either the design-implied grid intervals or the parameter space.

is positive for 43.36 percent of the subjects and negative for 56.64 percent. The right panel shows the corresponding distribution of individual posterior means of altruism in the advantageous domain (β_i). This individual posterior mean of β_i is positive for a large majority of subjects (86.81 percent) and negative for the remaining 13.19 percent.

To what extent are these two parameters interrelated? Figure 3 shows a two-dimensional binned scatter plot with a linear fit. The figure reveals that the two forms of altruism are *negatively correlated* (Spearman rank correlation: $\rho = -0.507$, $p < 0.01$): individuals who display higher degrees of altruism when ahead (a higher β) tend to show *lower* levels of altruism when behind (a lower α). While this negative correlation might be surprising at first, it is consistent with models of inequality aversion. For example, in the Fehr and Schmidt (1999) model (FS), individuals might be both averse to advantageous inequality (a higher β_{FS}) and averse to disadvantageous inequality (a higher α_{FS} in the original Fehr and Schmidt (1999) model).[11]

Figure 2: Distribution of individual posterior mean altruism parameters α_i and β_i



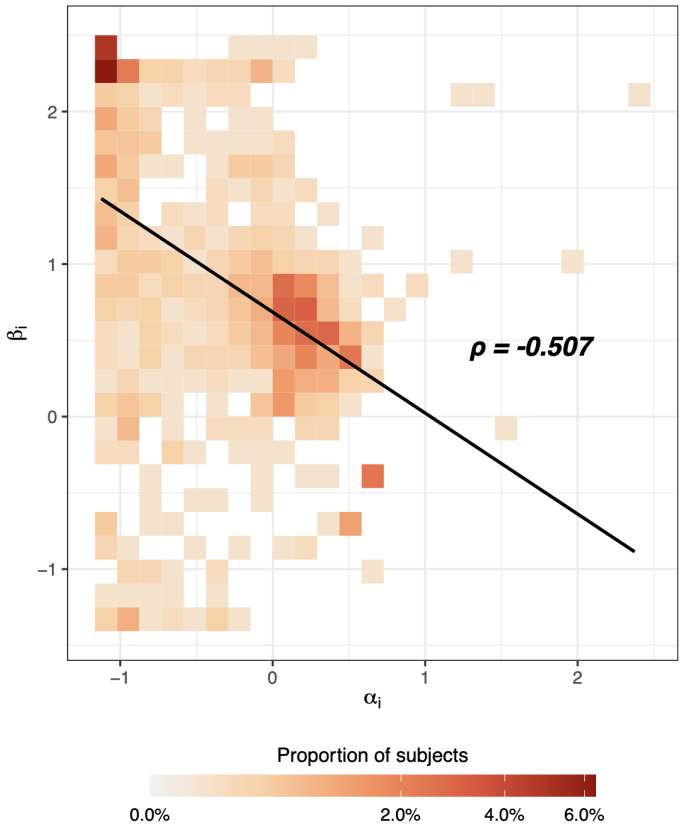
Note. Histograms show the relative proportion of subjects for the individual posterior mean altruism parameters. α_i denotes altruism when behind, and β_i denotes altruism when ahead.

5 Empirical strategy and omitted variable bias

Virtually all the papers discussed in the introduction have focused on the role of altruism in the advantageous domain—typically proxied using a single, standard dictator

[11]Recall that, in our specification, altruism in the disadvantageous domain (α) is defined as the opposite of aversion to disadvantageous inequality in the Fehr-Schmidt (FS) model, i.e., $\alpha \overset{def}{=} -\alpha_{\text{FS}}$.

Figure 3: Bivariate distribution of individual posterior mean altruism parameters (α_i, β_i)



Note. The figure shows a two-dimensional binned scatter plot. Fill color indicates the relative proportion of subjects in each bin. The solid line shows the fitted linear association between α_i and β_i. The Spearman rank correlation between individuals' mean posterior alpha and beta values is $\rho = -0.507$.

game (e.g., Kolstad and Lindkvist, 2013) or with unincentivized survey questions (e.g., Dohmen et al., 2009)—for labor market outcomes. Yet, we have shown above that the nature of respondents' social preferences is more complex: altruism in the advantageous and the disadvantageous domain coexist, and they are negatively correlated. What are the potential implications of neglecting these empirical regularities for the estimates of the effects of social preferences?

In Appendix C, we formally show that *omitting to control for one dimension of altruism leads to biased estimates of the effects of the other dimension of altruism*. The intuition behind this omitted-variable problem is straightforward. Suppose, for example, that when both dimensions are included in the regression, altruism in the advantageous domain and altruism in the disadvantageous domain both predict a stronger preference for working in a company that values social and environmental responsibility. Now consider a regression that includes only altruism in the disadvantageous domain as an explanatory variable. Because the two forms of altruism are negatively correlated, individuals who are more altruistic when behind also tend, on average, to be less altruistic when ahead. If altruism in the advantageous domain is omitted from the regression, this lower level of altruism when ahead is absorbed into the coefficient on altruism when behind. Since altruism in the advantageous domain is itself positively related to demand for CSR, omitting it induces a negative bias in the coefficient on altruism in the disadvantageous domain. As a result, the estimated coefficient on altruism in the disadvantageous domain is biased toward zero. In sufficiently strong cases, a truly positive and statistically significant relationship may therefore appear weak, insignificant, or even have the wrong sign when altruism in the advantageous domain is not controlled for.

Consequently, our analysis always follows the same structure. For each outcome variable, we estimate a "*full model*" that simultaneously accounts for both dimensions of altruism:

$$Z_i = \gamma_0 + \gamma_1 \alpha_i + \gamma_2 \beta_i + \gamma_3 X_i + \epsilon_i$$

where Z_i denotes the outcome variable of interest, α_i measures altruism in the disadvantageous domain, β_i measures altruism in the advantageous domain, X_i is a rich set of individual-level controls including age, gender, income, and education, and ϵ_i

is an error term.

In addition, we also report the estimates of two "*misspecified*" models that include only one dimension of altruism—either altruism when behind (α) or altruism when ahead (β). This exercise will help reveal the omitted-variable bias that arises if only one dimension of altruism is included in the regression. As we show below, failing to include both dimensions of altruism tends to yield estimates of social preferences that are biased towards zero, i.e., the role of social preferences will be *underestimated*.[12]

To facilitate interpretation and enable meaningful comparisons of effect sizes across specifications, we standardize α and β to have mean zero and standard deviation one. We also standardize the outcome variables, unless stated otherwise.

6 Results

We explore the role of social preferences in four steps. First, we examine how social preferences relate to individuals' life goals. Second, we investigate whether they predict respondents' preferred job-attributes. Third, we study whether social preferences are correlated with respondents' willingness to avoid working for the tobacco industry. Finally, we analyze the association between our experimentally elicited measures of social preferences and selection into jobs, exploiting a third-party dataset that records the occupation of a subset of our respondents.

6.1 The role of social preferences for life goals and aspirations

Result 1. *People with higher altruism in both the advantageous and the disadvantageous domain deem career-related aspects as less important for their overall life satisfaction.*

We first examine whether and how social preferences relate to individuals' life goals and aspirations. To that end, we analyze respondents' answers to the question "*Different people have different goals in life. Please indicate how important each of the following is to your overall life satisfaction.*" (Zukin and Szeltner, 2012). We are particularly

[12] In Table S.4, we summarize the effect of omitting α for the estimate of β (column 3), and the effect of omitting β for the estimate of α (column 4). The effects of social preferences are underestimated if the effects of both dimensions of altruism have the same sign in the full model (Case 1 and 2), and they are overestimated otherwise. As we will show, in our application omitting to include both dimensions of altruism leads to an underestimation of their effects.

interested in the answers given to the following four items, that relate to respondents' career and their financial well-being: *Having a prestigious career, Being financially secure, Being wealthy, Having a job where I can make an impact on causes or issues that are important to me.* We depict the results for these four career-related life goals in Table 2.[13] The table reveals several interesting insights.

First, people who display *higher* levels of altruism when behind rate having a prestigious career as being *less* important for their overall life satisfaction: a one-standard-deviation increase in α *reduces* the importance attached to career as an important driver of overall life satisfaction by 0.069 standard deviation ($p < 0.05$, column 1). Turning to altruism when ahead, a one-standard-deviation increase in β *reduces* the importance attached to career as an important driver of overall life satisfaction by 0.045 standard deviation, but the coefficient is insignificant in a regression that does not also control for α (column 2).

Second, we find that *both measures of altruism become quantitatively more important and gain statistical significance when we control for both of them simultaneously* (column 3). For example, the magnitude of the coefficient of altruism when ahead (β) almost doubles, from -0.045 to -0.089, when one simultaneously controls for altruism when behind, and the associated p-value decreases from more than 10 percent to less than 1 percent. In a similar vein, the magnitude of the coefficient associated with altruism when behind (α) increases from -0.069 to -0.106, and the significance increases from $p < 0.05$ to $p < 0.01$, when one controls for both measures simultaneously. This finding illustrates the omitted variable bias problem highlighted in Section 5. A regression that does not account for both dimensions of altruism will yield biased estimates because of the (negative) correlation between α and β. In our application, failing to control for both dimensions of altruism yields estimates that tend to underestimate the effects of social preferences.

In line with these results, columns 4-6 indicate that both forms of altruism are *negatively* associated with rating financial security as an important source of overall life satisfaction. Here too, we find that the magnitude and the statistical significance of both coefficients increase when moving from a specification that accounts for a

[13] In Figure S.3 in the Appendix, we depict the average importance attributed to each dimension for respondents' overall life satisfaction.

single form of altruism (columns 4 and 5) to a specification that includes both forms of altruism simultaneously (column 6). In addition, we find that less risk averse individuals (i.e., those with a higher willingness to take risks) attach a lower importance to financial security. Note, however, that the magnitude of the coefficients on altruism are even larger (in the advantageous case) or about the same (in the disadvantageous case) as the one on willingness to take risks.

Turning to the effects of wealth for overall life satisfaction (columns 7-9), we find that altruism when ahead is strongly negatively associated with deeming wealth as an important source of overall life satisfaction (columns 8 and 9, $p < 0.01$), but that altruism when behind appears to play no role. Interestingly, people who display stronger degrees of altruism do not deem having a job where *they* can make an impact as important for their overall life satisfaction (columns 10-12). This could be due to the fact that those with higher altruism do not value having more agency at work per se.[14] This interpretation is also consistent with Gomez-Martinez and Herz (2025), who measure intrinsic preferences for power in a large sample and find that strong preferences for power are associated with *weaker* preferences for equality.

In Table S.5 in the Appendix, we explore the relationship between altruism and the remaining life goals that are unrelated to career and wealth (*Having a partner*, *Having children*, and *Being a leader in my community*). We find some evidence that respondents with higher altruism, in particular altruism when ahead, find it less important to be a leader in their community. There is, however, no robust statistically significant association between either form of altruism and deeming family situation to be an important contributor to overall life satisfaction.

Together, these results indicate that people with stronger social preferences deem career and financial considerations less important to their overall life satisfaction. This raises the question whether social preferences play a role for the type of job they prefer and select into. We explore these questions in the next sections.

[14]Indeed, note that the survey item ("Having a job where I can make an impact on causes or issues that are important to me") is rather self-centered and is not related to other-regarding impacts.

Table 2: The relationship between social preferences and career-related life goals

	Career (z)			Fin. security (z)			Wealth (z)			Impact (z)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Altruism when behind	-0.069**		-0.106***	-0.042		-0.090***	0.029		-0.035	-0.007		-0.012
	(0.032)		(0.037)	(0.033)		(0.034)	(0.033)		(0.035)	(0.033)		(0.034)
Altruism when ahead		-0.045	-0.089**		-0.080**	-0.118***		-0.141***	-0.155***		-0.008	-0.013
		(0.031)	(0.036)		(0.032)	(0.033)		(0.033)	(0.037)		(0.033)	(0.034)
Female	-0.198***	-0.169**	-0.178**	0.051	0.085	0.078	-0.093	-0.055	-0.058	0.117*	0.121*	0.120*
	(0.070)	(0.071)	(0.070)	(0.070)	(0.069)	(0.069)	(0.072)	(0.072)	(0.072)	(0.070)	(0.071)	(0.071)
Age	-0.016***	-0.015***	-0.016***	-0.008**	-0.007**	-0.008**	-0.005	-0.004	-0.004	0.005	0.005	0.005
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Willingness to take risks (z)	0.062**	0.055*	0.052*	-0.075**	-0.085**	-0.088**	0.003	-0.013	-0.014	0.049	0.048	0.048
	(0.031)	(0.031)	(0.031)	(0.036)	(0.037)	(0.037)	(0.032)	(0.032)	(0.032)	(0.034)	(0.034)	(0.034)
Patience (z)	-0.009	-0.013	-0.008	-0.026	-0.029	-0.025	-0.033	-0.033	-0.032	-0.002	-0.002	-0.001
	(0.031)	(0.031)	(0.031)	(0.035)	(0.035)	(0.034)	(0.033)	(0.033)	(0.033)	(0.034)	(0.034)	(0.034)
Constant	0.788***	0.746***	0.779***	0.283	0.242	0.270	0.239	0.211	0.222	-0.139	-0.144	-0.140
	(0.186)	(0.187)	(0.186)	(0.191)	(0.188)	(0.189)	(0.187)	(0.186)	(0.187)	(0.182)	(0.182)	(0.182)
Income dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
H_0: $\alpha = \beta$			0.646			0.485			0.002			0.988
R^2	0.125	0.123	0.131	0.056	0.060	0.066	0.060	0.077	0.078	0.033	0.033	0.033
Observations	1030	1030	1030	1030	1030	1030	1030	1030	1030	1030	1030	1030

Notes: OLS regressions. The dependent variable is standardized (z-scored) and corresponds to respondents' rating of the importance of having a prestigious career (cols. 1-3), being financially secure (col. 4-6), being wealthy (col. 7-9), and having a job where they can make an impact (col. 10-12) for their overall life satisfaction. Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^*p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

6.2 The role of social preferences for preferences over work attributes

Result 2. *Higher altruism in the advantageous domain is associated with prioritizing CSR and meaning-at-work as important job-attributes. In the disadvantageous domain, higher altruism is associated with a stronger weight on CSR.*

In this section, we examine the role of social preferences for preferences over job-attributes, which we measured using the following question: "*Please look at the list below and tell us which one thing on this list you would most prefer in a job? Which comes next? Which is third-most important? etc.*", where respondents had to rank eight different job-attributes ranging from a high income to working for a company that values social and environmental responsibility, among others.

An advantage of eliciting a ranking is that it forces respondents to carefully consider each job-attribute and to make explicit tradeoffs between them. For example, the respondent cannot give the highest rating to both a high income and short working hours, but has to explicitly rank these attributes. For ease of interpretation, we recode the rank such that 8 means "most important" and 1 means "least important". We then use the rank of each item as a dependent variable. For transparency, we report the results of the regressions for each of the eight job-attributes.

We depict the average rating attributed to each of these job-attributes in Figure S.4 in the Appendix. On average, respondents rate financial security as the most important job-attribute (average rank = 6.13), and prospect of career advancement as the least important (average rank = 2.90).

In Table 3, we depict the regressions for two items that relate to corporate social responsibility (*Working for a company that values social and environmental responsibility*) and meaning (*Work is meaningful for (useful to) society*). On average, respondents rank these two job-attributes relatively high: Doing work that is meaningful is deemed the second most important job-attribute (average rating = 5.42) and working for a company that values CSR receives an average rating of 4.46. To what extent does altruism predict demand for these two job-attributes?

We find that respondents with higher levels of altruism rank working for a company that values CSR as significantly more important (see, e.g., column 3, $p < 0.01$

for both coefficients). The effect of altruism is quantitatively large. For example, a one-standard-deviation increase in altruism when ahead (β) leads to an increase in the importance attributed to working for a company that values CSR of 0.375, i.e., an increase of about 7 percent ($p < 0.01$). In comparison, female respondents rank working for a company that values CSR 0.345 points higher than male respondents (column 3, $p < 0.01$). Relatedly, we find that altruism when ahead significantly affects preferences for meaning-at-work: a one-standard-deviation increase in β yields an increase in the rating of meaning-at-work of 0.192 (Column 6, $p < 0.05$). However, altruism when behind has no significant effect.

How do respondents with stronger social preferences rank the other six job-attributes? In Table 4, we explore the effects of altruism for preferences over the different pecuniary aspects of a job (*High income, Job security, Chances for advancement*). We find that altruism *negatively* predicts preferences for the different pecuniary aspects of a job. For example, a one-standard-deviation increase in β leads to a decrease of 0.314 in the importance attributed to income (Column 3, $p < 0.01$). Altruism when behind, however, does not predict the weight attached to income. Relatedly, both dimensions of altruism negatively predict the importance attributed to job security and prospects of advancement (see, e.g., Columns 6 and 9, $p < 0.01$ on all coefficients for α and β). Together, these results are in line with those discussed in the previous section that people with stronger social preferences deem career-related aspects as less important for their overall life satisfaction.

Tables 3 and 4 provide further illustrations of the importance of including both dimensions of altruism in the regressions. For example, while altruism in the disadvantageous domain plays no role for CSR if one does not control for altruism in the advantageous domain (column 1 of Table 3), it plays a significantly positive role if one controls for the latter (column 3 of Table 3). Likewise, while altruism in the advantageous domain clearly reduces the importance of job security as a job attribute if both forms of altruism are included (column 6 in Table 4), it becomes insignificant if one fails to control for altruism in the disadvantageous domain (see column 5 in Table 4).

For completeness, we also explore the link between social preferences and preferences over the remaining job-attributes (*Short working hours, Flexibility in work hours,*

A work environment that allows to form friendships with colleagues) in Table S.6 in the Appendix. We don't find any systematic relationship between either form of altruism and preferences for short working hours, flexibility, or possibilities for making friendships at work, with one exception: those who are more altruistic when behind (higher α) also tend to rate short working hours as more important ($p < 0.01$).

Together, these results are aligned with the respondents' life goals that we analyzed in the previous section: those with higher altruism are *less* likely to deem career and financial situation to be important determinants of their overall life satisfaction. They are thus also *less* likely to rank the pecuniary aspects of a job as important. However, people with higher altruism attach significantly more importance to other aspects like working for a company that values CSR and doing work that is meaningful for society.

Table 3: The relationship between social preferences and meaning-at-work and corporate social responsibility (CSR)

	CSR			Meaning		
	(1)	(2)	(3)	(4)	(5)	(6)
Altruism when behind	0.040		0.193***	-0.027		0.051
	(0.066)		(0.074)	(0.071)		(0.078)
Altruism when ahead		0.295***	0.375***		0.170**	0.192**
		(0.068)	(0.076)		(0.070)	(0.077)
Female	0.429***	0.329**	0.345**	0.080	0.033	0.037
	(0.144)	(0.143)	(0.143)	(0.153)	(0.153)	(0.154)
Age	0.050***	0.047***	0.049***	0.041***	0.039***	0.040***
	(0.006)	(0.006)	(0.006)	(0.007)	(0.007)	(0.007)
Willingness to take risks (z)	0.009	0.043	0.049	-0.119*	-0.100	-0.098
	(0.067)	(0.067)	(0.067)	(0.069)	(0.068)	(0.069)
Patience (z)	0.137**	0.142**	0.133**	0.101	0.101	0.099
	(0.066)	(0.065)	(0.065)	(0.070)	(0.070)	(0.070)
Constant	2.376***	2.477***	2.416***	3.483***	3.519***	3.503***
	(0.396)	(0.396)	(0.395)	(0.420)	(0.419)	(0.420)
Income dummies	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes	Yes
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes
H_0: $\alpha = \beta$			0.022			0.095
R^2	0.155	0.171	0.177	0.092	0.097	0.097
Observations	1030	1030	1030	1030	1030	1030

Notes: OLS regressions. The dependent variable is the ranking given to working for a company that values social and environmental responsibility (col. 1-3) and doing work that is meaningful for society (col. 4-6). Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

Table 4: The relationship between social preferences and the pecuniary aspects of a job

	Income			Job security			Advancement		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Altruism when behind	0.133**		0.005	-0.162***		-0.232***	-0.118**		-0.213***
	(0.062)		(0.068)	(0.059)		(0.065)	(0.056)		(0.065)
Altruism when ahead		-0.316***	-0.314***		-0.074	-0.170***		-0.144**	-0.233***
		(0.061)	(0.067)		(0.057)	(0.061)		(0.060)	(0.067)
Female	-0.458***	-0.387***	-0.387***	-0.266*	-0.209	-0.228*	-0.220*	-0.150	-0.168
	(0.135)	(0.134)	(0.134)	(0.138)	(0.139)	(0.138)	(0.131)	(0.131)	(0.130)
Age	-0.038***	-0.037***	-0.037***	-0.001	0.002	0.000	-0.022***	-0.018***	-0.021***
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
Willingness to take risks (z)	0.034	0.000	0.000	-0.234***	-0.245***	-0.253***	0.207***	0.188***	0.182***
	(0.063)	(0.062)	(0.062)	(0.058)	(0.059)	(0.059)	(0.058)	(0.058)	(0.058)
Patience (z)	-0.015	-0.011	-0.012	0.039	0.030	0.041	0.026	0.018	0.028
	(0.061)	(0.060)	(0.060)	(0.062)	(0.062)	(0.062)	(0.058)	(0.058)	(0.058)
Constant	5.362***	5.330***	5.329***	6.218***	6.127***	6.200***	3.758***	3.667***	3.733***
	(0.359)	(0.356)	(0.357)	(0.350)	(0.352)	(0.351)	(0.334)	(0.338)	(0.333)
Income dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
H_0: $\alpha = \beta$			0.000			0.369			0.764
R^2	0.154	0.172	0.172	0.092	0.087	0.098	0.107	0.109	0.119
Observations	1030	1030	1030	1030	1030	1030	1030	1030	1030

Notes: OLS regressions. The dependent variable is the ranking given to a high income (col. 1-3), job security (col. 4-6), and chances of advancement (col. 7-9). Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

6.3 Are those with stronger social preferences willing to leave money on the table in order to avoid working for the tobacco industry?

Result 3. *People with higher altruism display a lower willingness to work for the tobacco industry.*

The tobacco industry tends to be perceived as immoral. As such, it is unlikely to be very attractive for people with stronger social preferences. However, the tobacco industry also tends to pay more than other industries (see, e.g., Schneider et al., 2024). This suggests that people with stronger social preferences might leave a significant amount of money on the table by selecting into alternative industries. In this section, we explore whether one can quantify the amount of money that those with stronger social preferences would be willing to give up in order to avoid working for the tobacco industry.

To answer this question, we analyze the results of a stylized discrete choice experiment in which respondents are asked to imagine that they have to choose between two companies that offered them a job: a large international tobacco company, and a medical research company. While the salary of the tobacco company is known (CHF 5,900 per month), the salary of the medical company is not yet known. Respondents are asked to reveal, for each possible salary offered by the medical company between CHF 5,000 and CHF 6,700, whether they would prefer working for the tobacco company or the medical company.[15]

For each respondent, we calculate the salary the respondent is willing to forgo to avoid working for the tobacco industry. A positive forgone salary means that the respondent is willing to give up salary in order to avoid working for the tobacco company, i.e., it indicates a preference *against* the tobacco company. In contrast, a negative forgone salary means that the respondent prefers to work for the tobacco company, even if it pays less than the medical company.[16] We depict the distribution

[15]See also Frank (1996), who elicits (hypothetical) compensating differential between different pairs of potential employers in a sample of Cornell undergraduates. For example, participants in his study report being willing to accept a pay-cut of about 50 percent to work as an ad copy writer for The American Cancer Society instead of doing the same job for Camel Cigarettes.

[16]The forgone salary ranges from CHF -900 (if the subject strongly prefers working for the tobacco company) to CHF + 900 (if the subject strongly prefers working for the medical company).

of forgone salaries in Figure S.5 in the Appendix E.1. The distribution is heavily skewed to the right, with 90.27% of the respondents willing to forgo some positive amount of money in order to avoid working for the tobacco company, and 47.28% not willing to work for the tobacco company even if it pays CHF 900 more than the medical company (a 15.25% higher salary). On average, respondents are willing to forgo CHF 519.78 in order to avoid working for the tobacco company.

Are people with stronger social preferences willing to give up more salary to avoid working for the tobacco company? To answer this question, we run a series of Tobit regressions that account for the fact that respondents' willingness to pay is censored both to the left and to the right. We report the results in Table 5. We find that those with stronger social preferences are willing to forego a substantial amount of salary in order to avoid working for the tobacco company. For example, a one-standard-deviation increase in α yields a CHF 54.37 increase in forgone salary (column 3, $p < 0.1$), which corresponds to 10.46 percent of the average willingness to pay in order to avoid working for the tobacco company. Relatedly, a one-standard-deviation increase in β yields a CHF 112.87 increase in forgone salary (column 3, $p < 0.01$)—21.71 percent of the average willingness to pay in order to avoid working for the tobacco company.

6.4 The role of social preferences for occupation choices

Result 4. *We find suggestive evidence that altruism in the advantageous domain is positively associated with employment in social jobs. This association is attenuated after accounting for respondents' preference for work that is meaningful for society, consistent with such preferences accounting for part of the relationship between altruism and occupational sorting.*

In the previous sections, we have shown that respondents with stronger social preferences attach less importance to the pecuniary aspects of a job and place greater weight on working for a company that values corporate social responsibility and on doing work that is meaningful for society. We have also shown that they are willing to forgo more salary to avoid working for the tobacco industry in a hypothetical choice experiment. We now ask whether these patterns extend beyond stated goals, ranked job attributes, and hypothetical choices to respondents' observed labor-market out-

Table 5: The relationship between social preferences and willingness to work for the tobacco industry

	Forgone salary (in CHF)		
	(1)	(2)	(3)
Altruism when behind	10.003		54.371*
	(27.726)		(28.878)
Altruism when ahead		91.605***	112.867***
		(30.031)	(31.047)
Female	257.897***	226.842***	230.750***
	(60.737)	(60.818)	(60.731)
Age	6.767**	5.856**	6.409**
	(2.765)	(2.737)	(2.761)
Willingness to take risks (z)	-11.204	0.691	2.667
	(29.152)	(29.549)	(29.509)
Patience (z)	17.527	18.751	16.378
	(28.560)	(28.071)	(28.120)
Constant	241.568	271.314*	252.420
	(158.941)	(158.366)	(158.959)
Income dummies	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes
Occupation dummies	Yes	Yes	Yes
H_0: $\alpha = \beta$			0.103
Pseudo R^2	0.010	0.011	0.011
Observations	1030	1030	1030

Notes: Tobit regressions, with dependent variable censored both to the left and to the right. The dependent variable is the forgone salary in CHF to avoid working for the tobacco company. Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^*p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

comes. To address this question, we use third-party data on respondents' occupations that were collected independently by our survey provider.

The occupation data distinguish 25 categories, including transportation, construction, arts and culture, education, health services, and social services. We classify education, health services, and social services as social jobs because the activities performed in these sectors are directly oriented toward providing benefits to other people. We define SocialJob as an indicator equal to one if a respondent is employed in one of these three sectors and zero otherwise. The analysis focuses on the 712 respondents who were working either full-time or part-time at the time of the study. Of these respondents, 138, or 19.38 percent, were employed in a social job.

Columns 1-3 of Table 6 reproduce our baseline specifications. When altruism in the advantageous domain (β) is entered on its own, its coefficient is 0.028 and statistically significant at the 10 percent level. When altruism in the advantageous and disadvantageous domains are included jointly, the coefficient on β is 0.024 and remains marginally significant. Because β is standardized, this estimate implies that a one-standard-deviation increase in altruism when ahead is associated with a 2.4 percentage-point increase in the probability of employment in a social job. Relative to the sample mean of 19.38 percent, this corresponds to an increase of approximately 12.4 percent. Altruism in the disadvantageous domain (α) is not significantly associated with employment in a social job, and women are about 13 percentage points more likely than men to work in these occupations.

The preceding results on job-attribute preferences suggest a possible explanation for the positive association between β and employment in social jobs. Section 6.2 shows that a one-standard-deviation increase in altruism in the advantageous domain raises the ranking of work that is meaningful for, or useful to, society by 0.192 positions ($p < 0.05$). This attribute captures an explicitly other-regarding component of a job: respondents may value such work not only because of its private benefits, but also because the work itself benefits other people. Education, health services, and social services are particularly likely to offer this type of job content. A preference for meaningful work may therefore be one way in which social preferences are related to occupational sorting.

Columns 4-6 examine this possibility by adding respondents' rankings of mean-

ingful work and corporate social responsibility. The preference for meaningful work is a strong predictor of employment in a social job. In column 4, a one-rank increase in the importance assigned to meaningful work is associated with a 3.9 percentage-point increase in the probability of working in a social job. Adding this variable reduces the coefficient on β from 0.024 to 0.016, a decline of one third, and the coefficient is no longer statistically significant. Its standard error remains unchanged at 0.014, indicating that the loss of significance reflects the smaller point estimate rather than a loss of precision. At the same time, R^2 rises from 0.118 to 0.165. By contrast, adding the ranking of corporate social responsibility alone leaves the coefficient on β unchanged at 0.024, while the corporate social responsibility ranking is itself unrelated to employment in a social job. When both job-attribute rankings are included in column 6, meaningful work remains strongly predictive, and the coefficient on β remains positive at 0.020 but is not statistically significant.

The baseline and augmented specifications answer different questions. Columns 1-3 estimate the overall, reduced-form association between altruism and employment in social jobs. Columns 4 and 6 ask whether altruism remains independently associated with social-job employment among respondents who attach the same importance to meaningful work. Importantly, the meaningful-work ranking is not necessarily a conventional confounder that should simply be controlled for. Under the preference-based account developed in this paper, it may itself be an intermediate outcome: social preferences may shape what individuals value in a job, and these job-attribute preferences may in turn be related to occupational sorting. Conditioning on meaningful work can therefore remove part of the very pathway through which altruism is associated with employment in social jobs.

Accordingly, the attenuation of the β coefficient should not be interpreted as showing that altruism is irrelevant for occupational sorting. Rather, it provides an additional insight into why altruism and social-job employment are related. The pattern is consistent with more altruistic individuals being more likely to choose socially useful occupations partly because they place greater value on the other-regarding content of the work. At the same time, the estimates do not establish a causal mediation mechanism. Respondents' occupations may have been chosen before either preference measure was collected, employment in a social job may itself strengthen

the stated preference for meaningful work, and some unobserved characteristics may affect all three variables. Moreover, the fact that a coefficient is statistically significant in one specification but not in another does not by itself establish that the two coefficients differ statistically. We therefore interpret the occupational-sorting evidence as suggestive.

Taken together, these results extend the predictive reach of our experimentally elicited measures of social preferences to an independently measured labor-market outcome. They also sharpen the paper's preference-based account of occupational sorting: altruism in the advantageous domain is positively associated with employment in social jobs, and a non-negligible part of this association appears to be linked to the greater importance that more altruistic individuals attach to work that is meaningful for society.

6.5 Discussion

Desire for consistency and demand effects. One potential concern in studies like ours is experimenter demand effects and desire for consistency (see, e.g., Haaland et al., 2023; Falk and Zimmermann, 2013). For example, a subject who behaves prosocially in the dictator game might want to appear consistent and indicate a preference for working in a company that values social and environmental responsibility if that question was asked directly after the dictator game. There are several reasons why we believe that this is very unlikely to be a problem in our setting. First, our staggered design, which decouples the measurement of social preferences from the life-goals block and other career-related preferences, largely mitigates this issue. Second, our evidence on sorting into social jobs relies on third-party data collected independently of our study and therefore cannot be subject to demand effects or consistency biases. Third, recent methodological work suggests that even strong experimentally-induced demand effects tend to result in only very modest behavioral responses (De Quidt et al., 2018). Moreover, we added a number of "filler" questions throughout our surveys in order to further mitigate such concerns.

Reverse Causality. Throughout the paper, we have interpreted our results as supporting a relationship that runs *from* social preferences *to* preferences over life goals,

Table 6: The relationship between social preferences and selection into social jobs

	Working in a social job					
	(1)	(2)	(3)	(4)	(5)	(6)
Altruism when behind	-0.020		-0.010	-0.016	-0.010	-0.014
	(0.015)		(0.015)	(0.014)	(0.014)	(0.014)
Altruism when ahead		0.028*	0.024*	0.016	0.024*	0.020
		(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
Female	0.135***	0.134***	0.132***	0.130***	0.132***	0.133***
	(0.033)	(0.033)	(0.033)	(0.032)	(0.033)	(0.032)
Age	0.001	0.001	0.001	-0.001	0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Willingness to take risks (z)	-0.020	-0.016	-0.017	-0.012	-0.017	-0.011
	(0.014)	(0.015)	(0.015)	(0.015)	(0.015)	(0.015)
Patience (z)	-0.002	-0.003	-0.002	-0.007	-0.002	-0.005
	(0.015)	(0.015)	(0.015)	(0.014)	(0.015)	(0.014)
Meaning				0.039***		0.042***
				(0.006)		(0.006)
CSR					-0.001	-0.013*
					(0.007)	(0.007)
Constant	0.022	0.024	0.027	-0.111	0.029	-0.091
	(0.078)	(0.077)	(0.077)	(0.077)	(0.080)	(0.079)
Income dummies	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes	Yes
H_0: $\alpha = \beta$			0.064	0.077	0.064	0.060
R^2	0.116	0.118	0.118	0.165	0.118	0.169
Observations	712	712	712	712	712	712

Notes: OLS regressions. The dependent variable is a dummy variable taking the value 1 if the respondent works in a social job. Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^*p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

job-attributes, and selection into different types of jobs, but could the causality run the other way? While it is of course possible that the social preferences of some of our subjects are shaped by their job and other life experiences, we believe that this is unlikely to be the driver of our findings for the following reasons. First, the social preferences of adults tend to be stable and unlikely to be strongly affected by short or medium term experiences. For example, we show in a companion paper that the social preferences of a general population *panel* sample are relatively stable between two measurement waves that occurred three years apart, both in the aggregate and at the individual level (Fehr ⓡ al., 2023). Second, if our effects were entirely driven by socialization at the workplace, one would expect much weaker or even no association between social preferences, life goals and career-related preferences among the youngest subjects in our sample, who have either not been exposed to the labor market yet, or only for a very short time. In the Appendix E.2, we explore the nature of the relationship between social preferences, life goals, and preferences over job-related aspects for the youngest 25 percent of our sample, consisting of respondents aged at most 35. For these young adults, the socialization channel is attenuated, i.e., the risk of reverse causality is mitigated. While our power to detect significant effects is weaker due to a much smaller sample size, we still find evidence that subjects with stronger social preferences attach less importance to monetary considerations for their overall life satisfaction, and that they are much more likely to rate working for a company that values CSR as a more important job-attribute, while at the same time putting less weight on the pecuniary aspects of a job. These results broadly align with those documented with the full sample. Finally, if the subjects who selected into social jobs became generally more prosocial, one would expect this more prosocial tendency to be reflected across the whole set of measures that we collected. Yet, we find that social preferences only predict job-related life goals. This context specificity of our findings is hard to reconcile with reverse causality. Altogether, we believe that it is very unlikely that socialization and preference endogeneity are the main drivers of our findings.

7 Conclusion

People differ in their aspirations and life goals, which may have implications for their labor supply. In this paper, we explored the role of people's social preferences in explaining these differences. Using a two-wave staggered study with a sample that is broadly representative of the adult Swiss population, we have shown that those with stronger social preferences are less likely to deem career and pecuniary aspects as important for their overall life satisfaction. Consistent with these life goals, more altruistic respondents rate the pecuniary aspects of a job and prospects of career advancement as relatively less important. However, they are much more likely to deem meaning-at-work and working for a company that values corporate social responsibility as very important. They are also more likely to be willing to forgo earnings in order to avoid working for the tobacco industry in a hypothetical choice experiment.

We also provided suggestive evidence that these preferences affect sorting in the labor market. Using third-party data on respondents' occupation that we can match to our study, we show that those who are more altruistic in the advantageous domain are more likely to be employed in a "social job" like education, health, and social services at the time of our study. We also document this relationship in the youngest quartile of respondents, which mitigates the risk of reverse causality through socialization at the workplace. Together, these results highlight the role of social preferences in providing a common account of people's life goals, their preferences over job-attributes, and their sorting into jobs with an other-regarding dimension.

In addition, our findings also have important implications for future research on the influence of social preferences. We show that altruism in the advantageous domain (β) and in the disadvantageous domain (α) are negatively correlated, and that failing to account for both forms of altruism yields biased estimates of the relationship between social preferences, life goals and preferences for job attributes. This follows from the omitted variable bias formula and is confirmed empirically: only including a single form of altruism produces biased estimates that tend, in our applications, to underestimate the role of social preferences. We believe that this result is important for the growing literature on external validity and the predictive ability of social preferences.

For organizations and practitioners, our findings highlight the importance of worker-organization fit along dimensions that extend beyond pay and career advancement. Organizations with a social mission may be particularly well positioned to attract workers with stronger social preferences by offering jobs whose contribution to society is visible and by credibly emphasizing their social and environmental commitments. Conversely, firms operating in sectors perceived as socially harmful may face a more limited pool of such workers or may have to offer compensating wage differentials as documented in Schneider et al. (2024). The attenuation of the association between altruism and employment in social jobs once preferences for meaningful work are taken into account provides an additional insight: social preferences may shape occupational sorting partly by shaping the job attributes workers value. Meaning and mission may therefore serve as matching attributes between workers and organizations, rather than merely as nonpecuniary benefits provided after hiring.

References

Alfaro, Laura, Ester Faia, Nora Lamersdorf, and Farzad Saidi (2022) "Health externalities and policy: The role of social preferences," *Management Science*, 68 (9), 6751–6761.

Barfort, Sebastian, Nikolaj A Harmon, Frederik Hjorth, and Asmus Leth Olsen (2019) "Sustaining honesty in public service: The role of selection," *American Economic Journal: Economic Policy*, 11 (4), 96–123.

Barr, Abigail and Pieter Serneels (2009) "Reciprocity in the workplace," *Experimental Economics*, 12 (1), 99–112.

Bauer, Rob, Tobias Ruof, and Paul Smeets (2021) "Get real! Individuals prefer more sustainable investments," *The Review of Financial Studies*, 34 (8), 3976–4043.

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.

Burbano, Vanessa C (2016) "Social responsibility messages and worker wage requirements: Field experimental evidence from online labor marketplaces," *Organization Science*, 27 (4), 1010–1028.

Burbano, Vanessa, Nicolas Padilla, and Stephan Meier (2024) "Gender differences in preferences for meaning at work," *American Economic Journal: Economic Policy*, 16 (3), 61–94.

Buurman, Margaretha, Josse Delfgaauw, Robert Dur, and Seth Van den Bossche (2012) "Public sector employees: Risk averse and altruistic?" *Journal of Economic Behavior & Organization*, 83 (3), 279–291.

Campos-Mercade, Pol, Claes Ek, Magnus Söderberg, and Florian Schneider (2025) "Social preferences and environmental externalities,"Technical report, Working Paper.

Campos-Mercade, Pol, Armando N Meier, Florian H Schneider, and Erik Wengström (2021) "Prosociality predicts health behaviors during the COVID-19 pandemic," *Journal of public economics*, 195, 104367.

Carpenter, Jeffrey and Erick Gong (2016) "Motivating agents: How much does the mission matter?" *Journal of Labor Economics*, 34 (1), 211–236.

Carpenter, Jeffrey and Caitlin Knowles Myers (2010) "Why volunteer? Evidence on the role of altruism, image, and incentives," *Journal of Public Economics*, 94 (11-12), 911–920.

Cassar, Lea (2019) "Job mission as a substitute for monetary incentives: Benefits and limits," *Management Science*, 65 (2), 896–912.

Cassar, Lea and Jesper Armouti-Hansen (2020) "Optimal contracting with endogenous project mission," *Journal of the European Economic Association*, 18 (5), 2647–2676.

Cassar, Lea and Stephan Meier (2018) "Nonmonetary incentives and the implications of work as a source of meaning," *Journal of Economic Perspectives*, 32 (3), 215–238.

Charness, Gary and Matthew Rabin (2002) "Understanding social preferences with simple tests," *Quarterly Journal of Economics*, 817–869.

Chater, Nick and George Loewenstein (2016) "The under-appreciated drive for sense-making," *Journal of Economic Behavior & Organization*, 126, 137–154.

Cohn, Alain and Michel André Maréchal (2018) "Laboratory measure of cheating predicts school misconduct," *The Economic Journal*, 128 (615), 2743–2754.

Cotofan, Maria, Lea Cassar, Robert Dur, and Stephan Meier (2023) "Macroeconomic conditions when young shape job preferences for life," *Review of Economics and Statistics*, 105 (2), 467–473.

Dai, Zhixin, Fabio Galeotti, and Marie Claire Villeval (2018) "Cheating in the lab predicts fraud in the field: An experiment in public transportation," *Management Science*, 64 (3), 1081–1100.

De Quidt, Jonathan, Johannes Haushofer, and Christopher Roth (2018) "Measuring and bounding experimenter demand," *American Economic Review*, 108 (11), 3266–3302.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde (2009) "Homo reciprocans: Survey evidence on behavioural outcomes," *The Economic Journal*, 119 (536), 592–612.

Dur, Robert and Robin Zoutenbier (2014) "Working for a good cause," *Public Administration Review*, 74 (2), 144–155.

Epper, Thomas F, Ernst Fehr, Claus Thustrup Kreiner, Søren Leth-Petersen, Isabel Skak Olufsen, and Peer Ebbesen Skov (2024) "Inequality aversion predicts support for public and private redistribution," *Proceedings of the National Academy of Sciences*, 121 (39), e2401445121.

Epper, Thomas F. and Ivan Mitrouchev (2025) "Measuring hearts and minds: A validated survey module on inequality aversion and altruism," *The Journal of Economic Inequality*, 23 (3), 777–808, 10.1007/s10888-025-09697-2.

Falk, Armin, Anke Becker, Thomas J Dohmen, David Huffman, and Uwe Sunde (2016) "The preference survey module: A validated instrument for measuring risk, time, and social preferences."

Falk, Armin and Florian Zimmermann (2013) "A taste for consistency and survey response behavior," *CESifo Economic Studies*, 59 (1), 181–193.

Fehr, Ernst ⓡ Thomas Epper ⓡ Julien Senn (2023) "The Fundamental Properties, the Stability, and the Predictive Power of Distributional Preferences," *URPP Equality of Opportunity Discussion Paper Series (45)*.

——— (2024) "Social preferences and redistributive politics," *Review of Economics and Statistics*, advance online publication.

Fehr, Ernst and K Schmidt (1999) "A Theory Of Fairness, Competition, and Cooperation," *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.

Frank, Robert H (1996) "What price the moral high ground?" *Southern Economic Journal*, 63 (1), 1–18.

Friebel, Guido, Michael Kosfeld, and Gerd Thielmann (2019) "Trust the police? Self-selection of motivated agents into the German police force," *American Economic Journal: Microeconomics*, 11 (4), 59–78.

Galizzi, Matteo M and Daniel Wiesen (2018) "Behavioral experiments in health economics."

Gartenberg, Claudine, Andrea Prat, and George Serafeim (2019) "Corporate purpose and financial performance," *Organization Science*, 30 (1), 1–18.

Gill, Andrej, Matthias Heinz, Heiner Schumacher, and Matthias Sutter (2023) "Social preferences of young professionals and the financial industry," *Management Science*, 69 (7), 3905–3919.

Gomez-Martinez, Francisco and Holger Herz (2025) "Who Wants Power? Measuring Intrinsic Preferences for Power and Their Behavioral Signature."

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart (2023) "Designing information provision experiments," *Journal of economic literature*, 61 (1), 3–40.

Hanna, Rema and Shing-Yi Wang (2017) "Dishonesty and selection into public service: Evidence from India," *American Economic Journal: Economic Policy*, 9 (3), 262–290.

Henderson, Rebecca and Eric Van den Steen (2015) "Why do firms have "purpose"? The firm's role as a carrier of identity and reputation," *American Economic Review*, 105 (5), 326–330.

Henkel, Aljosha ⓡ Ernst Fehr ⓡ Julien Senn ⓡ Thomas Epper (2025) "Beliefs about inequality and the nature of support for redistribution," *Journal of Public Economics*, 246, 105350.

Karlsson, Niklas, George Loewenstein, Jane McCafferty et al. (2004) "The economics of meaning," *Nordic Journal of Political Economy*, 30 (1), 61–75.

Kerschbamer, Rudolf and D Müller (2020) "Social preferences and political attitudes: An online experiment on a large heterogeneous sample," *Journal of Public Economics*, 182.

Kolstad, Julie Riise and Ida Lindkvist (2013) "Pro-social preferences and self-selection into the public health sector: evidence from an economic experiment," *Health policy and planning*, 28 (3), 320–327.

Kosse, Fabian and Michela M Tincani (2020) "Prosociality predicts labor market success around the world," *Nature communications*, 11 (1), 5298.

Lagarde, Mylene and Duane Blaauw (2014) "Pro-social preferences and self-selection into jobs: Evidence from South African nurses," *Journal of Economic Behavior & Organization*, 107, 136–152.

Li, Jing, Lawrence P Casalino, Raymond Fisman, Shachar Kariv, and Daniel Markovits (2022) "Experimental evidence of physician social preferences," *Proceedings of the National Academy of Sciences*, 119 (28), e2112726119.

Li, Jing, William H Dow, and Shachar Kariv (2017) "Social preferences of future physicians," *Proceedings of the National Academy of Sciences*, 114 (48), E10291–E10300.

Mas, Alexandre and Amanda Pallais (2017) "Valuing alternative work arrangements," *American Economic Review*, 107 (12), 3722–3759.

——— (2020) "Alternative work arrangements," *Annual Review of Economics*, 12 (1), 631–658.

Müller, Stephan and Holger A Rau (2021) "Economic preferences and compliance in the social stress test of the COVID-19 crisis," *Journal of Public Economics*, 194, 104322.

Non, Arjan, Ingrid Rohde, Andries De Grip, and Thomas Dohmen (2022) "Mission of the company, prosocial attitudes and job preferences: A discrete choice experiment," *Labour Economics*, 74, 102087.

Riedl, Arno and Paul Smeets (2017) "Why do investors hold socially responsible mutual funds?" *the Journal of Finance*, 72 (6), 2505–2550.

Rosso, Brent D, Kathryn H Dekas, and Amy Wrzesniewski (2010) "On the meaning of work: A theoretical integration and review," *Research in organizational behavior*, 30, 91–127.

Schneider, Florian H, Fanny Brun, and Roberto A Weber (2024) "Sorting and wage premiums in immoral work," *Review of Economics and Statistics*, 1–46.

Schneider, Florian H, Pol Campos-Mercade, Stephan Meier, Devin Pope, Erik Wengström, and Armando N Meier (2023) "Financial incentives for vaccination do not have negative unintended consequences," *Nature*, 613 (7944), 526–533.

Wiswall, Matthew and Basit Zafar (2018) "Preference for the workplace, investment in human capital, and gender," *The Quarterly Journal of Economics*, 133 (1), 457–507.

Wrzesniewski, Amy, Jane E Dutton, and Gelaye Debebe (2003) "Interpersonal sensemaking and the meaning of work," *Research in organizational behavior*, 25, 93–135.

Zukin, Cliff and Mark Szeltner (2012) "What workers want in 2012," *Talent Report for Rutgers University, https://www. issuelab. org/resources/15518/15518. pdf*.

ONLINE APPENDIX

A Additional material related to the survey waves and third party data S-2
A.1 Details on the measurement of covariates from wave 1 . . . . . S-2
A.2 Details on the measurement of covariates from wave 2 . . . . . S-2
A.2.1 Career and life goals block S-2
A.2.2 Additional measures S-4
A.3 Third party data on type of occupation S-4

B Demographic characteristics of sample population S-5
B.1 Descriptives S-5
B.2 Attrition S-7
B.3 Are there cross-study spillovers? S-9

C Why does neglecting to control for *both* α and β lead to biased estimates? S-10

D Structural estimation approach S-12

E Further Tables and Figures S-15
E.1 Choice experiment S-21
E.2 Effects on youngest respondents S-22

A Additional material related to the survey waves and third party data

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), financial situation, 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.

Additional measures included in wave 1

This first wave of data collection also included a host of survey items aimed at measuring respondents' beliefs. In particular, we measured beliefs about inequality, beliefs about social mobility, beliefs about poverty, beliefs about the determinants of success, trust in politicians. These items are unrelated to the present study. We report on them in a companion paper (Henkel ⓡ al., 2025).

A.2 Details on the measurement of covariates from wave 2

A.2.1 Career and life goals block

Life goals

Different people have different goals in life. Please indicate how important each of the following is to your overall life satisfaction. [1. Not at all important, 2. Rather not important, 3. Neither important nor unimportant, 4. Rather important, 5. Very important]

- Having a partner / being married
- Being wealthy
- Having children
- Having a prestigious career
- Being financially secure
- Being a leader in my community
- Having a job where I can make an impact on causes or issues that are important to me

Job attributes

Please look at the list below and tell us which one thing on this list you would most prefer in a job? Which comes next? Which is third-most important? etc. Please rank the following items in your preferred order, where 1 means "most important" and 8 means "least important". To move an item, click on it and drag it to the desired position ("drag and drop" function).

- High income
- Job Security
- Working hours are short, lots of free time
- Flexibility in work hours
- Chances for advancement
- Work is meaningful for (useful to) society
- A work environment that allows to form friendships with colleagues
- Working for a company that values social and environmental responsibility

Choice experiment (Tobacco industry)

Imagine that you are looking for a new job and that two companies offered you a job.

- Company A is a large international tobacco company.
- Company B is a medical research company which helps to develop new cures for diseases.

Both companies will give you some control over what you do and how you do it, but in both companies your boss will tend to set rather strict guidelines. Company A offers you a monthly salary of CHF 5,900. You are still unsure about the monthly salary that will be offered to you by company B. For each of the possible monthly salaries that company B might offer you, please tell us whether you prefer to take the job offer from company A or the job offer from company B.

1. Tobacco company for 5,900 □ | □ Medical company for 6,700
2. Tobacco company for 5,900 □ | □ Medical company for 6,600

   ...

17. Tobacco company for 5,900 □ | □ Medical company for 5,100
18. Tobacco company for 5,900 □ | □ Medical company for 5,000

A.2.2 Additional measures

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. (2016).

In addition, the second wave of data collection included a host of survey items aimed at measuring respondents' beliefs about inequality, and their preferences for redistribution. This wave also included an information intervention which we report on in a companion paper (Henkel ⓡ al., 2025): half of the participants were exposed to information aimed at correcting their beliefs about the extent of inequality in Switzerland. As we show in Appendix B.3 below, this information intervention did not affect any of the outcome variables that relate to respondents' preferences for life goals and job-attributes, which are at the center of this paper.

A.3 Third party data on type of occupation

We obtained a third-party dataset containing the type of occupation in which our respondents are employed. Occupations are divided into the following categories:

1. **Health**
2. **Social Services**
3. **Education**
4. **Health and Social Services**
5. Agriculture, Forest, Fishing
6. Industry
7. Construction
8. Energy
9. Retail (big)
10. Retail (small)
11. Retail/Repair (Cars, Machines, . . .)
12. Banking and Finance
13. Insurance
14. Transport
15. Communication
16. IT
17. Hotels/Tourism
18. Arts and Culture
19. Services (Economic, technical, scientific; e.g., consulting, research)
20. Other services (hair dresser, waste management, . . .)
21. Public Sector
22. Chemistry / Pharma
23. Sales / Customer Service / Office Work
24. Real Estate
25. Architecture
26. Don't know

In our analysis, we code "Education", "Health services", "Social services" and "Health and social services" as "SocialJob."

B Demographic characteristics of sample population

B.1 Descriptives

We depict the main descriptive statistics in Table S.1. Column 1 depicts the population mean, while column 2 documents the sample averages. Overall, our sample is broadly representative of the Swiss population with respect to age, gender, geographical area, and income. Table S.2 provides further details on our sample.

Table S.1: Descriptive statistics and balance checks

	Population	Sample
Age (mean)	51.1	47.55
Female	0.52	0.48
French-speaking	0.25	0.25
Income bracket : ≤ CHF 4000	0.28	0.32
Income bracket : CHF 4001-6000	0.26	0.22
Income bracket : CHF 6001-8000	0.22	0.20
Income bracket : CHF 8001-10000	0.12	0.11
Income bracket : CHF 10001-15000	0.09	0.07
Income bracket : ≥ CHF 15000	0.03	0.02
Income bracket : NA	-	0.05

Note: The table displays descriptive statistics of the Swiss population (column 1) and of our sample (column 2). The descriptive statistics include age (mean), the share of female respondents, 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).

Table S.2: Further descriptive statistics

	Mean	S.D.
Age	47.55	15.23
Female	0.48	0.50
Education: Obligatory school (up to 14 y.o.)	0.02	0.15
Education: Vocational training	0.37	0.48
Education: High school	0.11	0.31
Education: University	0.38	0.49
Education: Other	0.12	0.33
Income bracket: ≤ CHF 4000	0.32	0.47
Income bracket: CHF 4001-6000	0.22	0.42
Income bracket: CHF 6001-8000	0.20	0.40
Income bracket: CHF 8001-10000	0.11	0.32
Income bracket: CHF 10001-15000	0.07	0.25
Income bracket: > CHF 15000	0.02	0.15
Income bracket: NA	0.05	0.22
Occupation: Full-time worker	0.42	0.49
Occupation: Part-time worker	0.27	0.44
Occupation: Student	0.07	0.25
Occupation: Pensioner	0.17	0.37
Occupation: Unemployed	0.03	0.17
Occupation: Other	0.04	0.20
Observations	1031	

B.2 Attrition

1383 subjects participated in wave 1 and 1031 subjects completed wave 2, i.e., the rate of attrition is 25.5%. In Table S.3, we show that participation in the second wave is orthogonal to both altruism parameters (α and β) and to the bulk of the observable characteristics. Note that older respondents are slightly more likely *not* to drop out between waves, whereas students are slightly more likely to drop out.

Table S.3: Attrition

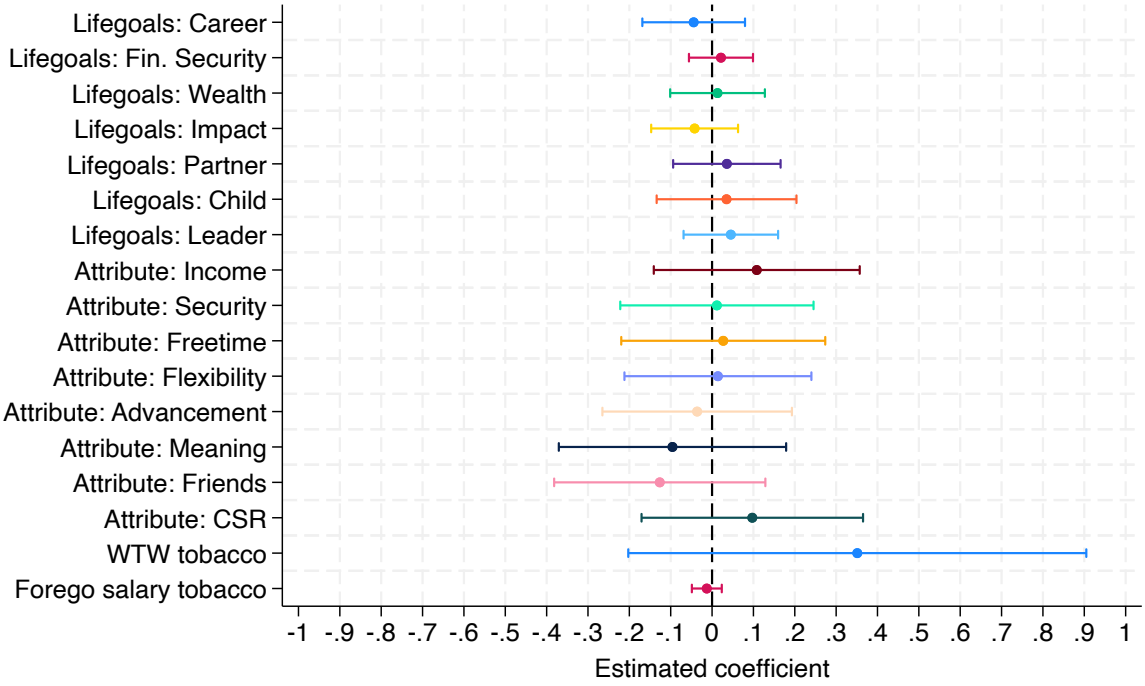
	(1)	(2)	(3)	(4)	(5)	(6)
alpha	-0.028					0.012
	(0.023)					(0.024)
beta	-0.012					-0.012
	(0.015)					(0.016)
Female		0.029				0.029
		(0.023)				(0.026)
Age		-0.005***				-0.005***
		(0.001)				(0.001)
French speaking		0.020				0.015
		(0.027)				(0.027)
Income bracket: ≤ CHF 4000			0.014			-0.012
			(0.057)			(0.057)
Income bracket: CHF 4001-6000			-0.026			-0.009
			(0.059)			(0.059)
Income bracket: CHF 6001-8000			-0.070			-0.038
			(0.059)			(0.060)
Income bracket: CHF 8001-10000			-0.009			0.036
			(0.064)			(0.065)
Income bracket: CHF 10001-15000			-0.041			0.005
			(0.069)			(0.072)
Income bracket: > CHF 15000			0.043			0.086
			(0.095)			(0.097)
Education: Obligatory school (up to 14 y.o.)				0.074		0.072
				(0.085)		(0.084)
Education: Vocational training				-0.020		-0.014
				(0.038)		(0.038)
Education: High school				0.101**		-0.004
				(0.049)		(0.050)
Education: University				0.068*		0.035
				(0.038)		(0.039)
Occupation: Full-time worker					0.052	0.025
					(0.059)	(0.064)
Occupation: Part-time worker					0.075	0.061
					(0.060)	(0.062)
Occupation: Student					0.246***	0.135*
					(0.071)	(0.078)
Occupation: Pensioner					-0.032	0.068
					(0.061)	(0.069)
Occupation: Unemployed					-0.085	-0.104
					(0.077)	(0.078)
Constant	0.273***	0.462***	0.271***	0.220***	0.196***	0.405***
	(0.017)	(0.041)	(0.053)	(0.033)	(0.056)	(0.098)
R^2	0.002	0.033	0.005	0.011	0.028	0.048
Observations	1383	1383	1383	1383	1383	1383

Note: OLS regression. The dependent variable is a dummy which equals 1 if the individual from the first wave did *not* participate in the second wave and equals 0 if the individual also participated in wave 2. Alpha and beta are the structurally estimated parameters. Income brackets are dummy variables that equal 1 if the respondent falls into the respective monthly income category. Education and Occupation dummies equal 1 if the respondent falls into the respective categories. Further socio-demographic variables include a dummy variable indicating whether the respondent is female, age, and a dummy variable indicating whether the respondent's native language is French. Robust standard errors in parentheses. Levels of significance: *$p < .1$, **$p < .05$, ***$p < .01$.

B.3 Are there cross-study spillovers?

As we discuss in Appendix A.2, the second wave of data collection contained an information provision experiment aimed at testing the effects of beliefs about inequality for support for redistribution. While beliefs about inequality have no direct connection with life goals and preferences over job-attributes, one might be worried that the information intervention accidentally affected the main outcome variables at the center of this paper. To investigate whether this is the case, we regressed each outcome variable used in this paper on a treatment dummy indicating whether the subjects were exposed to information about inequality in Switzerland.[17] We report the treatment effect (with standard errors) in Figure S.1 below. The figure unambiguously shows that the information intervention had *no* effect on our outcome variables, i.e., we can rule out that it affected any of the variables discussed in this paper.

Figure S.1: Effects of an information intervention aimed at shifting beliefs about income inequality on outcome variables used in this paper



[17] Only half of the subjects in Henkel ⓡ al. (2025) were exposed to the information intervention, the other half—in the control condition—were exposed to no information at all.

C Why does neglecting to control for *both* α and β lead to biased estimates?

Altruism in the advantageous domain and altruism in the disadvantageous domain are negatively correlated. What are the potential implications of neglecting this empirical regularity for the estimates of the effects of social preferences? In this Appendix, we illustrate in more detail that accounting only for altruism when ahead (β) and neglecting to simultaneously account for altruism when behind (α) leads to biased estimates of the effects of social preferences. For simplicity, we assume that the outcome variable is only affected by α and β.

Suppose, for example, that one is interested in estimating the effect of altruism when ahead on the variable Z and that one estimates the (misspecified) model

$$Z = b_0 + b_1\beta + u$$

when, in reality, the true model is

$$\begin{aligned} Z &= c_0 + c_1\beta + c_2\alpha + \varepsilon \\ E[\varepsilon \mid \beta, \alpha] &= 0. \end{aligned}$$

In the misspecified model, the effect of altruism when ahead is

$$\hat{b}_1 = \frac{\mathrm{Cov}(\beta, Z)}{\mathrm{Var}(\beta)}.$$

From the true model, we know that

$$\begin{aligned} \mathrm{Cov}(\beta, Z) &= \mathrm{Cov}\big(\beta, c_0 + c_1\beta + c_2\alpha + \varepsilon\big) \\ &= c_1 \mathrm{Var}(\beta) + c_2 \mathrm{Cov}(\beta, \alpha) + \underbrace{\mathrm{Cov}(\beta, \varepsilon)}_{=0} \\ &= c_1 \mathrm{Var}(\beta) + c_2 \mathrm{Cov}(\beta, \alpha). \end{aligned}$$

Thus

$$\hat{b}_1 = \frac{c_1 \mathrm{Var}(\beta) + c_2 \mathrm{Cov}(\beta, \alpha)}{\mathrm{Var}(\beta)} = c_1 + c_2 \frac{\mathrm{Cov}(\beta, \alpha)}{\mathrm{Var}(\beta)}.$$

That is, the bias in the misspecified regression is

$$\mathrm{Bias}(\hat{b}_1) = E[\hat{b}_1] - c_1 = c_2 \frac{\mathrm{Cov}(\beta, \alpha)}{\mathrm{Var}(\beta)} = \begin{cases} < 0 & \text{if } c_2 > 0 \\ > 0 & \text{if } c_2 < 0 \end{cases}$$

Because $\mathrm{Var}(\beta) > 0$, the sign of the bias is jointly determined by $\mathrm{Cov}(\beta, \alpha)$ and c_2. Since we have shown that altruism in the disadvantageous and advantageous domains are *negatively correlated* in our data (i.e., $\mathrm{Cov}(\beta, \alpha) < 0$), the sign of the bias of $\hat{b}_1$ is fully determined by the sign of c_2, the effect of altruism when behind on Z.

In Table S.4, we summarize the effect of omitting α for the estimate of β (column

3), and the effect of omitting β for the estimate of α (column 4). The effects of social preferences are underestimated if the effects of both dimensions of altruism have the same sign in the full model (Cases 1 and 2), and they are overestimated otherwise. As we show in the Results section, in our application omitting to control for both dimensions of altruism leads to an underestimation of their effects.

Table S.4: Implications of omitting to control for both dimensions of altruism for the estimates of α and β.

	Full model		Bias on $\hat{\beta}$ when α omitted	Bias on $\hat{\alpha}$ when β omitted
	Sign of c_2 (α)	Sign of c_1 (β)	Direction [Magnitude]	Direction [Magnitude]
Case 1	+	+	− (toward zero) [Underestimated]	− (toward zero) [Underestimated]
Case 2	−	−	+ (toward zero) [Underestimated]	+ (toward zero) [Underestimated]
Case 3	+	−	− (away from zero) [Overestimated]	+ (away from zero) [Overestimated]
Case 4	−	+	+ (away from zero) [Overestimated]	− (away from zero) [Overestimated]

Notes: Suppose that the true model is $Z = c_0 + c_1\beta + c_2\alpha + \varepsilon$, where Z is an outcome of interest, but that one omits to control for either α or β. One can show that $\text{Bias}(\hat{\alpha}) = c_1 \cdot \frac{\text{Cov}(\alpha,\beta)}{\text{Var}(\alpha)}$ and $\text{Bias}(\hat{\beta}) = c_2 \cdot \frac{\text{Cov}(\alpha,\beta)}{\text{Var}(\beta)}$. Since $\text{Cov}(\alpha, \beta) < 0$ in our data, the bias takes the opposite sign of the omitted coefficient. In Cases 1 and 2, both effects share the same sign, so omitting either parameter biases the remaining estimate toward zero, and the magnitude of the true effect is underestimated. In Cases 3 and 4, the effects have opposite signs, so the bias pushes the estimate away from zero, and the magnitude of the true effect is overestimated.

D Structural estimation approach

This appendix describes the estimation of the individual altruism parameters α_i and β_i and the choice-error scale λ_i. The notation $U_i(x_i, x_j)$, $\theta_i = (\alpha_i, \beta_i, \lambda_i)$, and A follows Section 4.2 in the main text. We index choice situations with $t = 1, ..., T$ and use A_t to refer to the alternatives presented in situation t. $a = (x_i, x_j) \in A_t$ refers to an allocation.

For the stochastic choice rule, we assume additive errors that are independent and identically distributed (i.i.d.) according to a standardized Type-I extreme value (EV1) distribution. Latent utility for allocation $a = (x_i, x_j) \in A_t$ is

$$\tilde{U}_{it}(a) = U_i(a) + \lambda_i \varepsilon_{it}(a), \quad \varepsilon_{it}(a) \overset{\text{i.i.d.}}{\sim} \text{EV1}(0, 1).$$

This error specification yields the multinomial logit choice probability $P_i(a \mid A_t, \theta_i)$ specified in Section 4.2. As described there, the choice-error scale $\lambda_i > 0$ governs how closely choices follow utility maximization.

Conditional on θ_i, observations for a given subject are treated as independent. Let a_{it}^{obs} denote subject i's observed allocation in choice situation t, and let $D_i = \left\{(A_t, a_{it}^{\text{obs}})\right\}_{t=1}^{T}$ denote that subject's choice data. For a candidate value θ of θ_i, the subject-level log likelihood is

$$\ell_i(\theta) = \sum_{t=1}^{T} \log P_i(a_{it}^{\text{obs}} \mid A_t, \theta).$$

Let $\mathcal{A}$, $\mathcal{B}$, and $\mathcal{L}$ denote the grids for α_i, β_i, and λ_i, respectively. Their Cartesian product $\Theta = \mathcal{A} \times \mathcal{B} \times \mathcal{L} = \{\theta_g\}_{g=1}^{G}$ is a finite joint grid, where $\theta_g = (\alpha_g, \beta_g, \lambda_g)$ and $G = |\Theta|$ is the number of grid points. Our main specification uses a discrete uniform prior over these points,

$$\pi_g = \Pr(\theta_i = \theta_g) = \frac{1}{G}, \quad g = 1, ..., G. \tag{3}$$

For subject i, the posterior mass assigned to grid point θ_g is

$$w_{ig} = \Pr(\theta_i = \theta_g \mid D_i) = \frac{\exp(\ell_i(\theta_g))\pi_g}{\sum_{h=1}^{G} \exp(\ell_i(\theta_h))\pi_h}.$$

Under the uniform prior in Equation 3, this reduces to normalized likelihood weights,

$$w_{ig} = \frac{\exp(\ell_i(\theta_g))}{\sum_{h=1}^{G} \exp(\ell_i(\theta_h))}.$$

Thus the maximum a posteriori (MAP) grid point coincides with the maximum likelihood estimate (MLE) constrained to the same grid. The posterior means we report and use in our analyses use the full distribution of normalized likelihood mass across the grid.

The α_i and β_i grids use deterministic cut-points implied by the experimental design. Each decision situation (see Figure 1a) implies a critical value of α_i or β_i, respectively. Graphically, these critical values are exactly where the (piecewise linear) in-

difference curves overlap the allocation lines (see also Epper and Mitrouchev (2025)). For a non-vertical allocation line, let

$$m = \frac{\Delta x_j}{\Delta x_i} = \frac{x_{j,2} - x_{j,1}}{x_{i,2} - x_{i,1}}$$

denote its slope. Our utility function gives

$$\frac{dU_i}{dx_i} = 1 + \alpha_i(m-1) \quad \text{for } x_i < x_j, \qquad \frac{dU_i}{dx_i} = 1 + \beta_i(m-1) \quad \text{for } x_i \geq x_j.$$

Setting the relevant derivative to zero gives $\alpha^* = \frac{1}{1-m}$ in the disadvantageous domain and $\beta^* = \frac{1}{1-m}$ in the advantageous domain. We set the independently specified λ_i grid to be log-spaced.

Our baseline prior assigns equal mass to every joint grid point. The α_i and β_i nodes have unequal spacing because they are derived from design cut-points, so equal node masses imply unequal prior mass per unit of continuous parameter space.

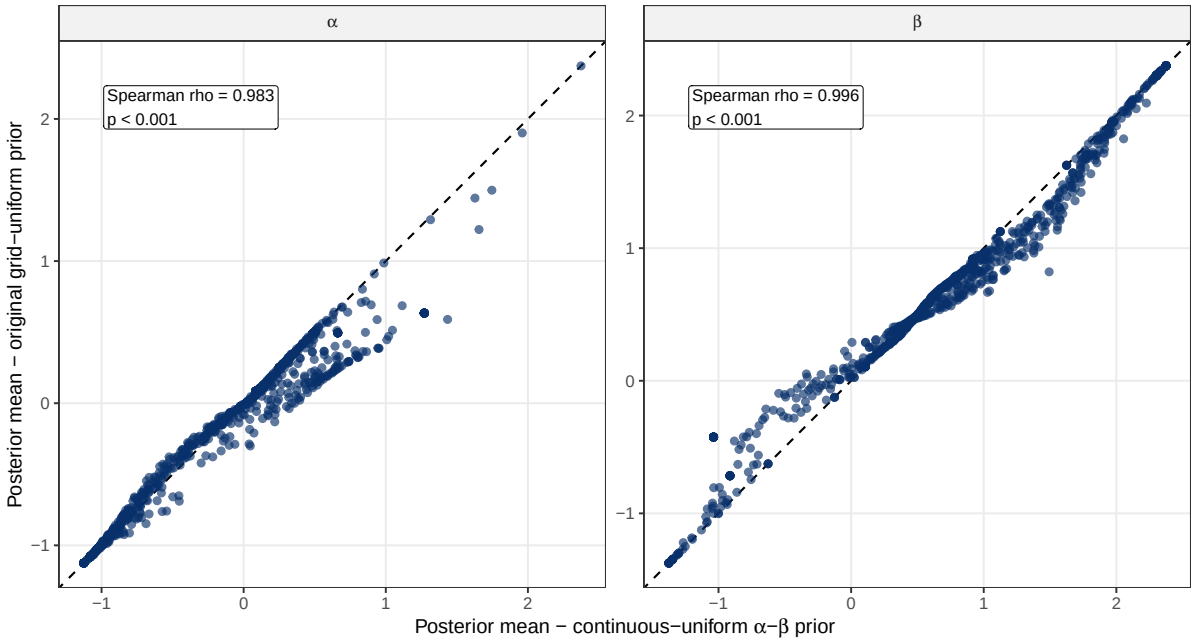
As a robustness check, we examine whether our structural estimates are sensitive to how the uniform prior is implemented. Our baseline specification assigns equal prior probability to each point on the parameter grid. Because the grid points are not equally spaced, this implies that different regions of the continuous parameter space receive different prior probabilities. We therefore re-estimate the model using an alternative prior that assigns equal probability density across the continuous parameter space. This alternative specification gives greater weight to grid points representing larger intervals, while leaving the likelihood function, parameter grid, and other estimation assumptions unchanged. Let $\Delta\alpha_a$ and $\Delta\beta_b$ denote the widths of the design-defined cells represented by nodes α_a and β_b. Midpoint quadrature assigns prior mass proportional to cell areas, and the prior assumptions for λ_i remain the same. Our prior then becomes

$$\pi^{\text{CU}}_{abl} \propto \Delta\alpha_a \Delta\beta_b, \quad w^{\text{CU}}_{iabl} = \frac{\exp(\ell_i(\alpha_a, \beta_b, \lambda_l))\Delta\alpha_a\Delta\beta_b}{\sum_{a',b',l'} \exp(\ell_i(\alpha_{a'}, \beta_{b'}, \lambda_{l'}))\Delta\alpha_{a'}\Delta\beta_{b'}}.$$

This robustness exercise uses the same likelihood, grid nodes, parameter support, and λ_i prior as the baseline. The sole change is therefore the prior weights over the α_i and β_i cells.

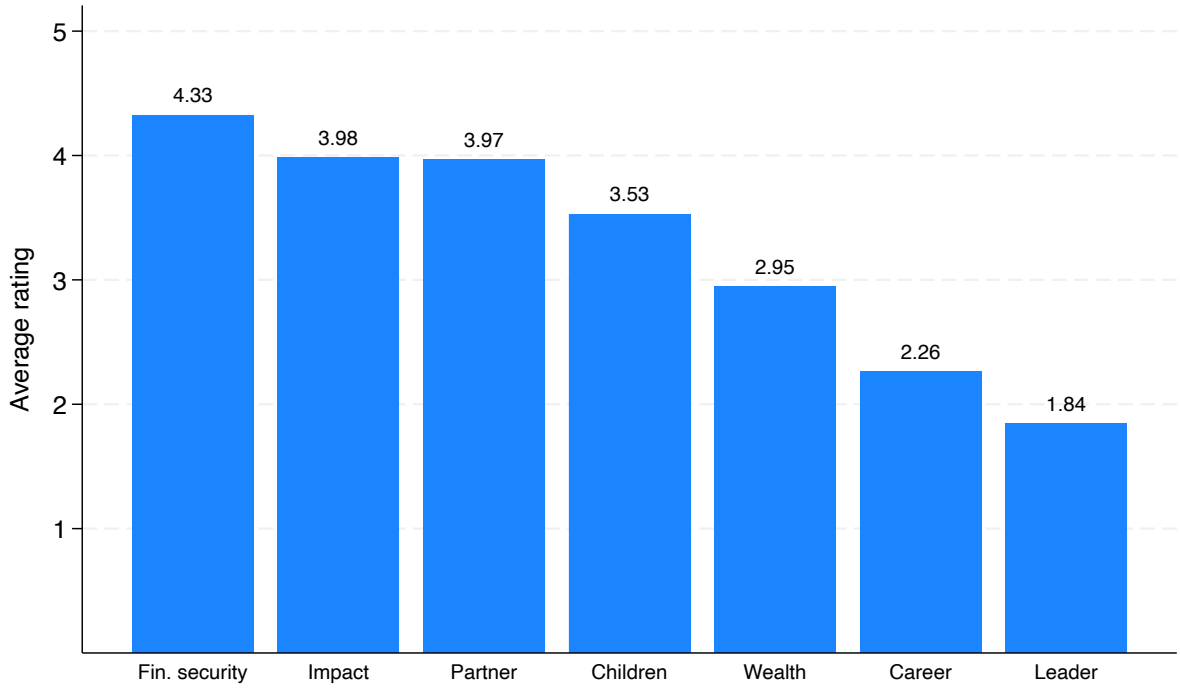
Figure S.2 compares the individual posterior mean estimates under our baseline specification that assigns equal prior probability to each point on the parameter grid (y-axis) to the alternative specification that assigns equal probability density across the continuous parameter space (x-axis). As the panels illustrate, the parameter estimates are only slightly perturbed by changes in the prior assumption. The Spearman rank-correlations between individual altruism parameters in the two domains are 0.983 and 0.996, respectively.

Figure S.2: Robustness of preference parameter estimates to alternative specifications of the uniform prior



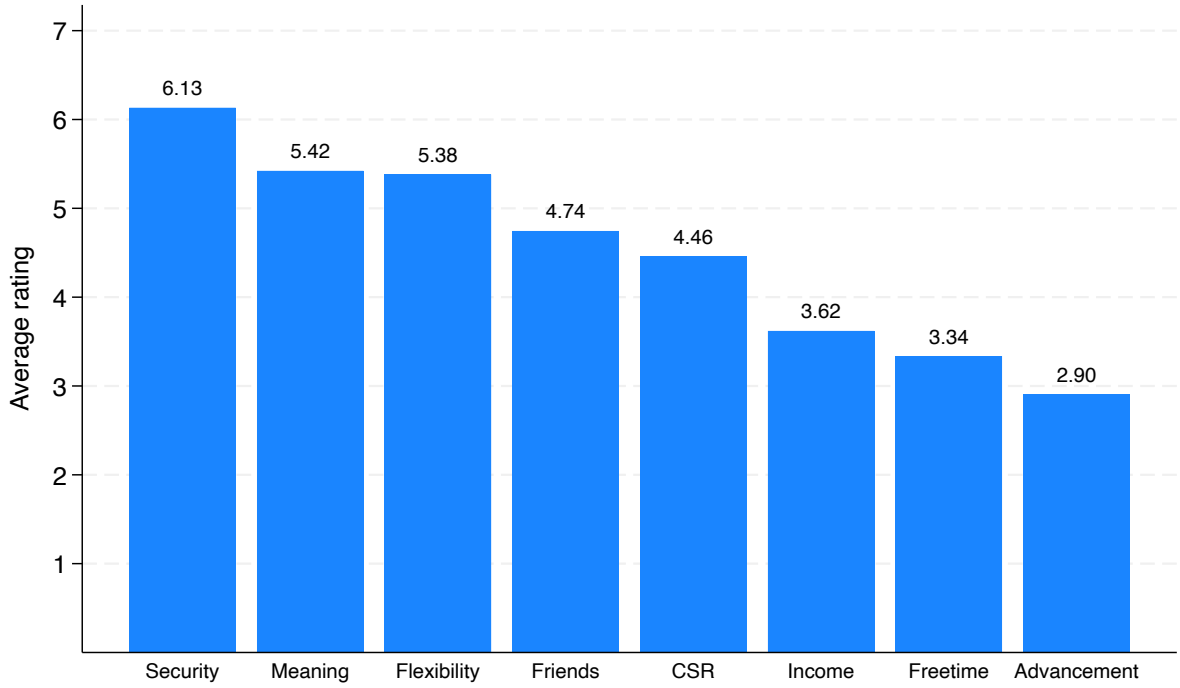
E Further Tables and Figures

Figure S.3: Life goals (raw data)



Notes: For each item, subjects were asked to answer the following questions "Different people have different goals in life. Please indicate how important each of the following is to your overall life satisfaction." on a scale from 1 (not at all important) to 5 (very important.)

Figure S.4: job-attributes (raw data)



Notes: Subjects were asked to rank the items from their most preferred to their least preferred ("Please look at the list below and tell us which one thing on this list you would most prefer in a job? Which comes next? Which is third-most important, etc. ?") The Figure depicts the average rating of each item, where higher values indicate *more preferred* job-attributes.

Table S.5: The relationship between social preferences and life goals that are unrelated to professional life

	Leader (z)			Partner (z)			Children (z)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Altruism when behind	-0.031		-0.063*	-0.063**		-0.057	-0.027		-0.028
	(0.032)		(0.037)	(0.032)		(0.035)	(0.033)		(0.035)
Altruism when ahead		-0.053*	-0.080**		0.037	0.013		0.009	-0.003
		(0.032)	(0.037)		(0.033)	(0.036)		(0.034)	(0.036)
Female	-0.270***	-0.247***	-0.252***	-0.233***	-0.232***	-0.236***	0.021	0.024	0.022
	(0.068)	(0.068)	(0.067)	(0.074)	(0.074)	(0.074)	(0.072)	(0.072)	(0.072)
Age	0.001	0.002	0.002	-0.002	-0.001	-0.002	0.004	0.004	0.004
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Willingness to take risks (z)	0.090***	0.083***	0.081**	-0.027	-0.024	-0.026	-0.001	-0.000	-0.001
	(0.032)	(0.032)	(0.032)	(0.033)	(0.033)	(0.033)	(0.034)	(0.035)	(0.035)
Patience (z)	-0.009	-0.012	-0.009	0.009	0.006	0.009	-0.005	-0.006	-0.005
	(0.031)	(0.031)	(0.031)	(0.035)	(0.035)	(0.035)	(0.033)	(0.033)	(0.033)
Constant	0.176	0.147	0.167	-0.061	-0.078	-0.060	-0.105	-0.114	-0.105
	(0.191)	(0.191)	(0.191)	(0.192)	(0.192)	(0.192)	(0.186)	(0.186)	(0.186)
Income dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ho: alpha=beta			0.652			0.065			0.527
R^2	0.074	0.076	0.079	0.036	0.034	0.036	0.030	0.029	0.030
Observations	1030	1030	1030	1030	1030	1030	1030	1030	1030

Notes: OLS regressions. The dependent variable is standardized (z-scored) and corresponds to respondents' rating of the importance of being a leader in their community (col. 1-3), having a partner (col. 4-6), and having children (col. 7-9) for their overall life satisfaction. Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

Table S.6: The relationship between social preferences and the non-pecuniary aspects of a job

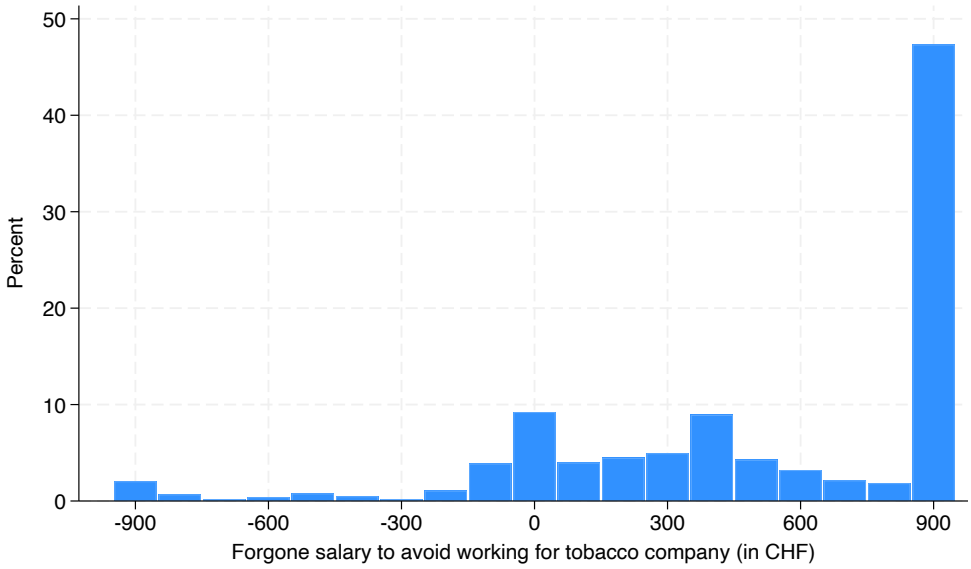
	Free time			Flexibility			Colleagues		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Altruism when behind	0.171***		0.199***	-0.021		-0.015	-0.016		0.011
	(0.065)		(0.071)	(0.058)		(0.063)	(0.065)		(0.070)
Altruism when ahead		-0.015	0.067		0.022	0.016		0.062	0.066
		(0.062)	(0.067)		(0.063)	(0.068)		(0.068)	(0.074)
Female	-0.095	-0.126	-0.110	0.133	0.131	0.130	0.396***	0.380**	0.381**
	(0.144)	(0.145)	(0.144)	(0.131)	(0.132)	(0.132)	(0.147)	(0.148)	(0.148)
Age	-0.029***	-0.032***	-0.030***	-0.011*	-0.011*	-0.011*	0.010	0.010	0.010
	(0.007)	(0.007)	(0.007)	(0.006)	(0.006)	(0.006)	(0.007)	(0.006)	(0.007)
Willingness to take risks (z)	0.004	0.005	0.011	0.017	0.019	0.019	0.082	0.089	0.090
	(0.066)	(0.066)	(0.066)	(0.062)	(0.062)	(0.062)	(0.069)	(0.070)	(0.070)
Patience (z)	-0.047	-0.039	-0.048	-0.071	-0.072	-0.071	-0.170**	-0.170**	-0.171**
	(0.062)	(0.061)	(0.062)	(0.059)	(0.058)	(0.059)	(0.067)	(0.066)	(0.066)
Constant	4.460***	4.529***	4.467***	5.896***	5.893***	5.897***	4.448***	4.458***	4.455***
	(0.369)	(0.368)	(0.369)	(0.346)	(0.347)	(0.347)	(0.388)	(0.388)	(0.389)
Income dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ho: alpha=beta			0.087			0.668			0.493
R^2	0.111	0.104	0.112	0.034	0.034	0.034	0.047	0.048	0.048
Observations	1030	1030	1030	1030	1030	1030	1030	1030	1030

Notes: OLS regressions. The dependent variable is the ranking given to having short working hours (col. 1-3), flexibility in working hours (col. 4-6) and having a work environment that allows to form friendships with colleagues (col. 7-9). Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^*p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

E.1 Choice experiment

In Figure S.5, we depict the distribution of respondents' willingness to pay to avoid working for the tobacco company.

Figure S.5: Willingness to give up salary in order to avoid working for the tobacco company



E.2 Effects on youngest respondents

In this Appendix, we explore the nature of the relationship between social preferences, life goals, and preferences over job-related aspects for the youngest 25 percent of our sample, consisting of respondents aged at most 35. For these young adults, the risk of reverse causality is strongly mitigated, as we argue in Section 6.5. While our power to detect significant effects is weaker due to a much smaller sample size, we still find evidence that subjects with stronger social preferences attach less importance to monetary considerations for their overall life satisfaction (Table S.7), and that they are strongly more likely to rate working for a company that values CSR as a more important job-attribute (Table S.8), while at the same time putting less weight on the pecuniary aspects of a job (Table S.9). These results align with those documented with the full sample.

Table S.7: The relationship between social preferences and career-related life goals (young respondents)

	Career (z)			Fin. security (z)			Wealth (z)			Impact (z)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Altruism when behind	-0.147**		-0.223**	-0.097		-0.193***	0.034		-0.057	-0.064		-0.061
	(0.071)		(0.086)	(0.070)		(0.069)	(0.068)		(0.076)	(0.068)		(0.073)
Altruism when ahead		-0.037	-0.148*		-0.090	-0.185***		-0.147**	-0.176**		0.037	0.006
		(0.068)	(0.085)		(0.062)	(0.061)		(0.063)	(0.071)		(0.066)	(0.071)
Female	-0.239*	-0.175	-0.210	-0.029	0.039	0.008	-0.250*	-0.207	-0.216*	0.288**	0.297**	0.287**
	(0.132)	(0.137)	(0.133)	(0.130)	(0.129)	(0.127)	(0.130)	(0.128)	(0.129)	(0.117)	(0.123)	(0.120)
Age	-0.019	-0.019	-0.017	-0.028*	-0.026	-0.025	-0.011	-0.009	-0.009	-0.002	-0.002	-0.002
	(0.017)	(0.018)	(0.017)	(0.017)	(0.017)	(0.016)	(0.017)	(0.017)	(0.017)	(0.016)	(0.017)	(0.017)
Willingness to take risks (z)	0.089	0.067	0.082	-0.060	-0.081	-0.068	-0.037	-0.049	-0.045	0.049	0.045	0.049
	(0.066)	(0.068)	(0.065)	(0.070)	(0.068)	(0.068)	(0.065)	(0.066)	(0.066)	(0.068)	(0.069)	(0.069)
Patience (z)	0.075	0.063	0.078	0.119*	0.111	0.123*	0.054	0.054	0.057	-0.028	-0.032	-0.028
	(0.064)	(0.065)	(0.063)	(0.070)	(0.070)	(0.068)	(0.065)	(0.064)	(0.064)	(0.066)	(0.067)	(0.066)
Constant	0.597	0.687	0.566	0.493	0.560	0.455	0.659	0.654	0.623	-0.751	-0.716	-0.749
	(0.558)	(0.560)	(0.550)	(0.558)	(0.548)	(0.547)	(0.492)	(0.490)	(0.490)	(0.539)	(0.538)	(0.542)
Income dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ho: alpha=beta			0.340			0.921			0.125			0.400
R^2	0.107	0.090	0.122	0.109	0.108	0.131	0.111	0.129	0.131	0.124	0.122	0.124
Observations	248	248	248	248	248	248	248	248	248	248	248	248

Notes: OLS regressions. The dependent variable is standardized (z-scored) and corresponds to respondents' rating of the importance of having a prestigious career (col. 1-3), being financially secure (col. 4-6), being wealthy (col. 7-9), and having a job where they can make an impact (col. 10-12) for their overall life satisfaction. Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^{*}p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

Table S.8: The relationship between social preferences and meaning-at-work and corporate social responsibility (young respondents)

	CSR			Meaning		
	(1)	(2)	(3)	(4)	(5)	(6)
Altruism when behind	0.305**		0.518***	0.034		0.196
	(0.141)		(0.149)	(0.164)		(0.178)
Altruism when ahead		0.156	0.413***		0.217	0.314*
		(0.145)	(0.143)		(0.166)	(0.181)
Female	0.756***	0.592**	0.676**	0.037	-0.056	-0.024
	(0.274)	(0.278)	(0.273)	(0.331)	(0.330)	(0.333)
Age	0.009	0.006	0.003	0.009	0.006	0.005
	(0.037)	(0.037)	(0.037)	(0.043)	(0.043)	(0.043)
Willingness to take risks (z)	0.226	0.280*	0.245*	-0.226	-0.198	-0.211
	(0.144)	(0.148)	(0.146)	(0.165)	(0.163)	(0.165)
Patience (z)	-0.103	-0.078	-0.112	0.150	0.156	0.143
	(0.139)	(0.138)	(0.138)	(0.169)	(0.167)	(0.170)
Constant	3.448***	3.253***	3.533***	3.615***	3.574***	3.680***
	(1.212)	(1.251)	(1.240)	(1.271)	(1.248)	(1.260)
Income dummies	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes	Yes
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes
Ho: alpha=beta			0.536			0.547
R^2	0.156	0.144	0.180	0.119	0.126	0.130
Observations	248	248	248	248	248	248

Notes: OLS regressions. The dependent variable is the ranking given to working for a company that values social and environmental responsibility (col. 1-3) and doing work that is meaningful for society (col. 4-6). Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^*p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

Table S.9: The relationship between social preferences and the pecuniary aspects of a job (young respondents)

	Income			Job security			Advancement		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Altruism when behind	-0.120		-0.347**	-0.373***		-0.387***	0.050		-0.072
	(0.146)		(0.172)	(0.126)		(0.140)	(0.138)		(0.162)
Altruism when ahead		-0.269**	-0.441***		0.165	-0.027		-0.201	-0.236
		(0.131)	(0.152)		(0.135)	(0.146)		(0.140)	(0.167)
Female	-0.522*	-0.380	-0.436	-0.042	0.025	-0.037	-0.259	-0.202	-0.213
	(0.279)	(0.278)	(0.276)	(0.268)	(0.273)	(0.269)	(0.293)	(0.286)	(0.289)
Age	-0.018	-0.014	-0.012	0.012	0.010	0.013	-0.061*	-0.058*	-0.058*
	(0.038)	(0.037)	(0.037)	(0.033)	(0.033)	(0.033)	(0.035)	(0.035)	(0.035)
Willingness to take risks (z)	0.200	0.155	0.179	-0.336**	-0.363***	-0.338**	0.119	0.103	0.108
	(0.133)	(0.134)	(0.132)	(0.132)	(0.136)	(0.133)	(0.148)	(0.150)	(0.151)
Patience (z)	0.173	0.160	0.183	0.061	0.036	0.062	0.348**	0.348**	0.352**
	(0.148)	(0.144)	(0.144)	(0.143)	(0.145)	(0.143)	(0.151)	(0.150)	(0.150)
Constant	5.054***	5.152***	4.963***	5.545***	5.749***	5.539***	4.560***	4.551***	4.512***
	(1.227)	(1.170)	(1.183)	(1.070)	(1.054)	(1.074)	(1.160)	(1.164)	(1.170)
Income dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ho: alpha=beta			0.556			0.018			0.301
R^2	0.099	0.110	0.127	0.105	0.083	0.105	0.137	0.144	0.145
Observations	248	248	248	248	248	248	248	248	248

Notes: OLS regressions. The dependent variable is the ranking given to a high income (col. 1-3), job security (col. 4-6), and chances of advancement (col. 7-9). Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^*p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.

Table S.10: The relationship between social preferences and life goals that are unrelated to professional life (young respondents)

	Leader (z)			Partner (z)			Children (z)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Altruism when behind	-0.005		-0.080	-0.149**		-0.148*	-0.177**		-0.128
	(0.071)		(0.082)	(0.069)		(0.086)	(0.069)		(0.084)
Altruism when ahead		-0.105	-0.145*		0.075	0.002		0.159**	0.096
		(0.067)	(0.077)		(0.072)	(0.086)		(0.077)	(0.092)
Female	-0.329**	-0.288**	-0.300**	-0.090	-0.066	-0.090	-0.104	-0.102	-0.122
	(0.127)	(0.124)	(0.125)	(0.136)	(0.137)	(0.136)	(0.136)	(0.134)	(0.134)
Age	-0.008	-0.006	-0.006	-0.031*	-0.032*	-0.031*	-0.023	-0.025	-0.024
	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)	(0.019)	(0.018)	(0.019)
Willingness to take risks (z)	0.162**	0.150**	0.155**	0.029	0.019	0.029	0.035	0.031	0.040
	(0.063)	(0.061)	(0.061)	(0.063)	(0.063)	(0.063)	(0.078)	(0.075)	(0.076)
Patience (z)	0.100	0.098	0.103	-0.015	-0.025	-0.015	-0.016	-0.026	-0.018
	(0.065)	(0.065)	(0.065)	(0.071)	(0.072)	(0.071)	(0.074)	(0.073)	(0.073)
Constant	0.737	0.751	0.707	0.512	0.593	0.512	0.419	0.508	0.439
	(0.565)	(0.568)	(0.570)	(0.568)	(0.573)	(0.568)	(0.590)	(0.585)	(0.588)
Income dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Education dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ho: alpha=beta			0.411			0.055			0.007
R^2	0.151	0.161	0.165	0.063	0.049	0.063	0.102	0.098	0.108
Observations	248	248	248	248	248	248	248	248	248

Notes: OLS regressions. The dependent variable is standardized (z-scored) and corresponds to respondents' rating of the importance of being a leader in their community (col. 1-3), having a partner (col. 4-6), and having children (col. 7-9) for their overall life satisfaction. Altruism when behind is measured by the standardized value of α. Altruism when ahead is measured by the standardized value of β. Robust standard errors in parentheses. Levels of significance: $^*p < 0.1, ^{**}p < 0.05, ^{***}p < 0.01$.