



HAL
open science

Empirical research on ethical preferences: how popular is prioritarianism?

Erik Schokkaert, Benoît Tarrow

► **To cite this version:**

Erik Schokkaert, Benoît Tarrow. Empirical research on ethical preferences: how popular is prioritarianism?. 2021. halshs-03110312

HAL Id: halshs-03110312

<https://shs.hal.science/halshs-03110312>

Preprint submitted on 14 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

WP 2104 – January 2021

Empirical research on ethical preferences: how popular is prioritarianism?

Erik Schokkaert, Benoît Tarrow

Abstract:

We survey the empirical literature on ethical preferences, covering both survey studies and incentivized laboratory experiments. Crucial axioms such as the Pigou-Dalton transfer principle are not accepted by a large fraction of the subjects. Moreover, in formulating their distributive preferences subjects attach much importance to the sources of income differences. Their preferences behind a veil of ignorance do not coincide with their preferences in the position of a social planner. These results suggest that prioritarian policy proposals will not necessarily be supported by a majority of the population. Although the majority opinion does not necessarily reflect the ethically desirable perspective, the empirical results still raise some interesting normative challenges.

Keywords:

Surveys, lab experiments, distributive preferences, prioritarianism, inequality aversion

JEL codes:

D63, D71

Empirical research on ethical preferences: how popular is prioritarianism?¹

Erik Schokkaert², Benoît Tarrow³

January 12, 2021

We survey the empirical literature on ethical preferences, covering both survey studies and incentivized laboratory experiments. Crucial axioms such as the Pigou-Dalton transfer principle are not accepted by a large fraction of the subjects. Moreover, in formulating their distributive preferences subjects attach much importance to the sources of income differences. Their preferences behind a veil of ignorance do not coincide with their preferences in the position of a social planner. These results suggest that prioritarian policy proposals will not necessarily be supported by a majority of the population. Although the majority opinion does not necessarily reflect the ethically desirable perspective, the empirical results still raise some interesting normative challenges.

Keywords: surveys, lab experiments, distributive preferences, prioritarianism, inequality aversion.

JEL classification: D63, D71.

¹This paper is written for the book *Prioritarianism in Practice*, edited by Matthew Adler and Ole Norheim, that will be published by Cambridge University Press. We thank the Editors and participants to presentations at the University of Bergen for helpful comments.

²Department of Economics, KU Leuven.

³Univ Lyon, University Lumière Lyon 2, GATE UMR 5824, F-69130 Ecully, France

1 Introduction

There are two strands in the literature on distributive justice. The first strand is empirical and aims at getting a better understanding of the feelings and attitudes in society with respect to distributive justice and how these are linked to the individual and social characteristics of the respondents. The second one is normative and discusses which arguments are valid for defining a “just” situation or a “just” society. It is possible to keep these two approaches completely separate. Ethical theories are usually more subtle and complex than everyday opinions and a survey that would be based on the former runs the danger of being far removed from any relevant psychological reality of lay citizens. Moreover, a specific perspective on distributive justice does not become ethically acceptable, just because it is supported by a majority of the population.

Yet, in a political democracy it is nearly impossible to implement any theory of justice without sufficient support from the general public. This support will depend on the citizen’s own values and preferences. As an example, if (as seems to be the case) for a majority of citizens the shape of the tax function matters independently of the resulting final outcomes, a prioritarian policy maker will have a hard time putting some of his specific policy proposals into practice. Even if one considers the majority opinions to be ethically unacceptable, one still has to convince a sufficient number of citizens if one wants to implement one’s own supposedly superior conception of justice. A better understanding of the empirical structure of the uninformed opinions, that have “to be corrected“, may then be extremely useful, even necessary.

One could even go further and also accept that knowledge of the empirical results can contribute to ethical reasoning.¹ Following the Rawlsian model of the reflective equilibrium, the challenge of normative thinking is to construe a consistent view of the world in which “relevant normative facts” concerning specific matters in the real world cohere with general principles. The former should perhaps not be restricted to the own intuitions of the ethical observer as in the traditional Rawlsian position. In fact, these intuitions may be strongly biased if that ethical observer is a white, highly educated, male, economics, philosophy or law professor with an income in the upper decile of the distribution. The quest for a reflective equilibrium becomes more challenging and relevant if we include among the relevant normative facts also the “judgments” of other people, i.e. the “common-sense” opinions in society (Miller, 1992) - and even real “justice related” behavior (Güth and Kliemt, 2010).

In this chapter we look at prioritarianism from this empirical perspective - and with the double purpose of getting a better insight into the societal feasibility of prioritarian policies and of collecting relevant findings that may challenge traditional ethical thinking. We first discuss some results that are derived from questionnaire studies. It would be inappropriate to derive any behavioral predictions from these studies, as the incentives in the real world are very different from those in a questionnaire environment. Actual behavior will normally not be driven primarily by normative

¹See Gaertner and Schokkaert, 2012, for a more elaborate justification of this position

considerations, but it will reflect a mixture of self-interest, norms, and the desire to make a good impression on other people. Yet, the aim of this research is to derive information about subjective norms and ethical convictions, and not to predict behavior. In a study about norms, we do not want the subjects to be guided by self-interested considerations, and questionnaire studies without monetary payments are then a natural approach. We will argue that these methods can (and perhaps should) exploit the possibilities of quasi-experimental designs. In section 2 we summarize some relevant results that are obtained in large survey studies with representative samples of the population, while section 3 is devoted to more structured quasi-experimental questionnaire studies. It will turn out that it is not easy to really test for the popular acceptance of prioritarianism, but that it still can be safely concluded that some of its main ingredients (such as the Pigou-Dalton transfer principle) are put in question by respondents. We also give an overview of estimates of the degree of inequality aversion and show that the results are sensitive to the framing of the questions. Some more basic issues about the acceptance of the veil of ignorance as an intellectual device and about the traditional welfarist framework are discussed in sections 3.4 and 3.5 respectively.

Section 4 summarizes the results from real behavioral experiments in which subjects are financially incentivized. These experiments are useful if we want to understand real behavior. Yet, the section also shows that cleverly designed experiments yield indeed normatively relevant facts. To get insights into ethical preferences, rather than e.g. into the degree of altruism or envy, we focus on experimental settings in which one subject takes the position of a social planner or in which subjects have to decide behind a veil of ignorance. As with the questionnaires, there are not many studies that aim directly at checking the relevance of prioritarianism, but we will see that there are relevant results on the trade-off between equity and efficiency and, more specifically, on the acceptance of the (crucial) transfer axiom. We will also discuss some findings on the choice between ex ante and ex post approaches in a context of risk.

In section 5 we come back to the issue of the subtle relationship between the empirical work surveyed in this chapter and the normative debate about prioritarianism. Section 6 is a brief conclusion.

We focus in this chapter on questionnaire and experimental studies of the attitudes of the population. Other related questions are not treated. We do not discuss inverse optimum studies, that try to derive the underlying value judgments of decision makers from real world policy decisions (some of this work is summarized in chapter 4, this volume). We do not elaborate on the link between inequality and subjective well-being (see Clark and d'Ambrosio, 2015, for a survey on this topic). Finally, we focus on inequality aversion and not on risk aversion. There is a substantial literature on the similarities and differences between these two notions, and studies have investigated in some detail whether decisions under risk and redistributive decisions follow the same logic. We will comment on the issue when we discuss the results concerning the veil of ignorance, but a complete summary of this literature goes beyond the boundaries of this chapter.

2 Ethical preferences as measured in large surveys

Questionnaire studies on redistributive preferences or attitudes have always played a prominent role in the work of sociologists and political scientists, but in recent years they have also become more popular among economists. The resulting literature is huge and it is impossible to summarize it in this brief contribution. We will therefore limit ourselves to some basic findings from the economic literature that seem most relevant for prioritarianism. Yet we will argue that, while these findings are highly relevant to understand which redistributive policies are politically feasible in the real world, they do not really give a deeper insight into the specific normative questions relating to prioritarianism. For that purpose more focused approaches are necessary. These will be discussed in section 3.

2.1 Many people want redistribution, but there are large interindividual differences in distributive preferences

Most often, the dependent variable in the survey studies is derived from the answers on simple questions such as:

“Some people think that the government in Washington should do everything to improve the standard of living of all poor Americans (they are at point 1 on this card). Other people think it is not the government’s responsibility and that each person should take care of himself (they are at point 5). Where are you placing yourself in this scale?” (General Social Survey - GSS)

“It is the responsibility of the government to reduce the difference in income between people with high incomes to those with low incomes” (International Social Survey Program - ISSP)

“The government should take measures to reduce differences in income levels” (European Social Survey - ESS)

Figure 1 (taken from Kuziemko et al., 2015) gives a broad picture of the answers on these questions for the US. The results for other countries are sometimes slightly different, but the overall picture is rather robust. A majority of the population is in favor of redistribution, but (not surprisingly) support is more outspoken among low income earners. However, support for redistribution did not increase in the last decades, despite the fact that there is near consensus that income inequality in the US did strongly increase in that period. We first comment on the first finding and then on the second (in section 2.2).

The general picture in Figure 1 hides a lot of interindividual variation. In their survey chapter, Alesina and Giuliano (2011) point to the importance of self-interest. This captures the direct effect of own income (as in the Figure), but also the indirect effect of inequality on own income, e.g. through the increase in criminal behavior. Personal experiences also play a role. These include, for instance, misfortune in the past (growing up in a recession increases the willingness to redistribute, Giuliano and Spilimbergo, 2014), or the number of siblings and the position of the respondent in the order of children (younger children are more in favor of redistribution, Yamamura, 2015).

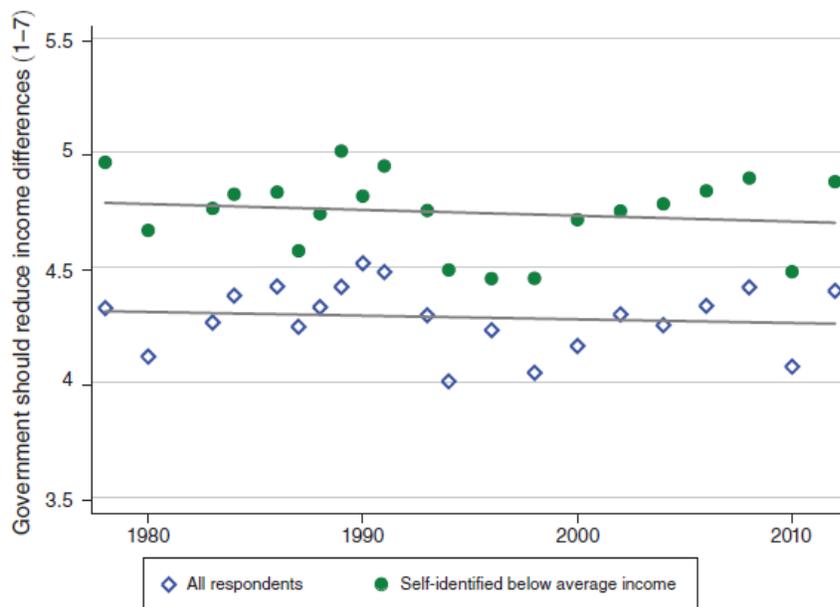


Figure 1: “Government should reduce income differences” (GSS) - average responses in the US over the last decades (Source: Kuziemko et al., 2015)

A second important factor is the belief in social mobility, which has become known as the “prospect of upward mobility” (POUM)-hypothesis (Benabou and Ok, 2001). This phenomenon has been further analyzed in Alesina et al. (2018). Respondents who believe that there is a substantial degree of social mobility in society, are less in favor of income redistribution, as they believe that everybody (including themselves) has a fair chance to improve his own position. Perceptions of mobility are unrealistic, however. Americans overestimate the degree of mobility in their society, while Europeans overestimate the risk of getting trapped in poverty. We will come back to the importance of (mis)perception, but the empirical results on mobility clearly suggest that respondents take a life-time perspective when evaluating distributions.

An essential finding in the literature is the overwhelming importance attached by respondents to the causes of income differences. At least since the work of Fong (2001), it has become generally accepted that respondents make a distinction between factors for which individuals should be held responsible (their “effort”) and factors for which they should be compensated (“luck”). We will come back to the normative implications of this finding in section 3.5, but it is certainly fair to state that there is consensus about the descriptive relevance of the distinction. This does not mean that everybody gives the same content to what is “effort” and what is “luck”, or has the same estimate about the relative importance of effort and luck in the determination of incomes. Different opinions or perceptions in this regard seem to be among the main factors underlying differences in ethical preferences, e.g. between males and females or between different educational or ethnic

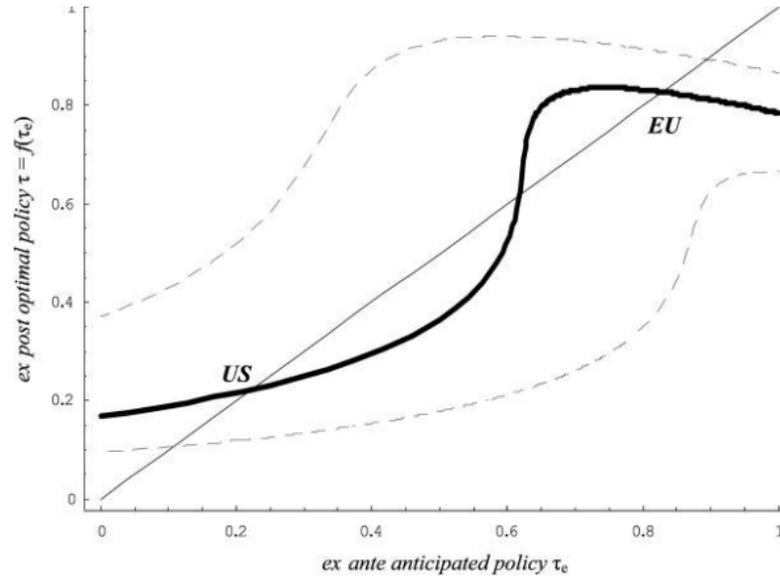


Figure 2: Preferences and policies: multiple equilibria (Alesina and Angeletos, 2005)

groups (Schokkaert and Truyts, 2017).

In fact, Alesina and Angeletos (2005) point out that the interaction between ethical preferences and policy outcomes may be at the heart of intercountry differences in both. They show that there is a significant positive relationship at the cross-country level between the percentage of respondents who believe that income is mainly determined by luck rather than by effort, and social spending as a percentage of GDP. This positive relationship does not necessarily imply that attitudes “cause” policies. Causality also runs in the opposite direction, with the percentage of social spending influencing both the justice perceptions of individuals and the link between luck and income in the real world. This interaction can result in multiple equilibria, as illustrated in Figure 2. In the “American” equilibrium, low taxes make effort very important in the determination of incomes, inducing ethical preferences supporting low taxes. The “European” equilibrium is characterized by a high level of redistribution, undermining the importance of effort in income formation. Whether one accepts the specific model of Alesina and Angeletos or not², it seems clear that attitudes of respondents are strongly influenced by the institutions in which they grow up. Intercultural differences in preferences even persist after people have moved to another country (Luttmer and Singhal, 2011).

As we have seen, perceptions of mobility or of the relative importance of effort and luck in

²In the Alesina-Angeletos model, perceptions of the relative importance of effort and luck are realistic. Benabou and Tirole (2006) have an alternative model which gives a much greater role to misperceptions and biases in explaining the differences between Europe and the US. Moreover, Isaksson and Lindskog (2009) present some empirical evidence to question the starting point that different weights given to effort and luck would be the main driving force of intercountry differences in redistributive preferences.

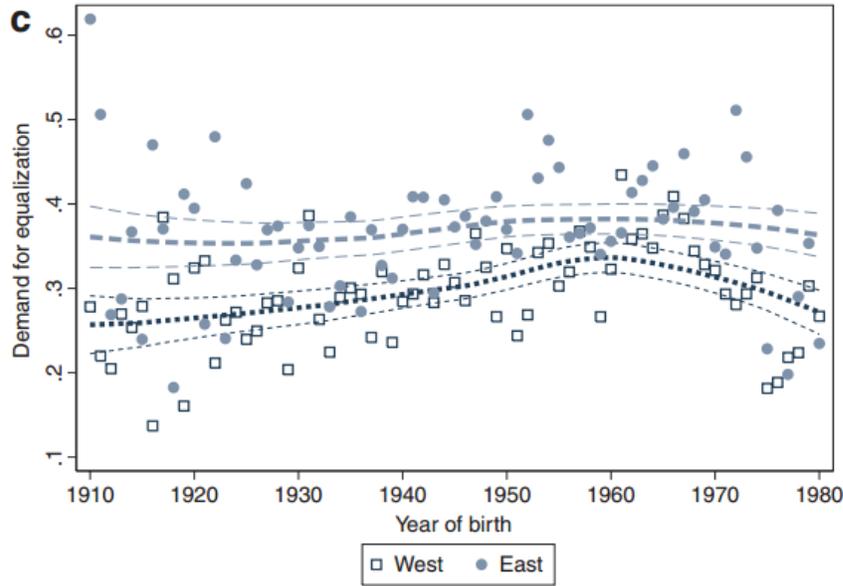


Figure 3: Willingness to redistribute by Eastern and Western Germans (source: Kuhn 2013)

the determination of incomes, are crucial to interpret correctly the answers of the respondents. In general, neglecting the basic (and almost trivial) insight that respondents start from their perception of the actual situation when evaluating the desirability of redistribution, may lead to completely mistaken conclusions about their true ethical preferences. A striking illustration is given by the work of Kuhn (2013) on the differences in the willingness to redistribute of Eastern and Western Germans. Kuhn uses the ISSP-questions (of the waves 1992 and 1999) on the “estimated” and “ethically justified” wages of different occupations to calculate Gini inequality measures. Figure 3 shows the degree of redistribution, i.e. the reduction in the Gini, that is desired by his respondents. The results suggest that Eastern Germans are more in favor of redistribution at all possible ages. It is tempting to relate this finding to the fact that Eastern Germans have been influenced by the values of communism.

Yet, this conclusion should be interpreted cautiously, as illustrated in Figure 4. Panel a of that Figure shows the inequality as perceived by Eastern and Western Germans. The former clearly perceive more inequality. Panel b describes the desired inequality: there are no longer any differences between the two groups (again, whatever the age of the respondents)! The seemingly stronger willingness to redistribute, suggested by Figure 3, does not reflect any difference in inequality aversion, but only differences in the perception of actual wage differentials. Note that very many studies do not collect this information on perceived inequality, and just ask about the willingness to redistribute (see the formulations of the questions at the beginning of this section).

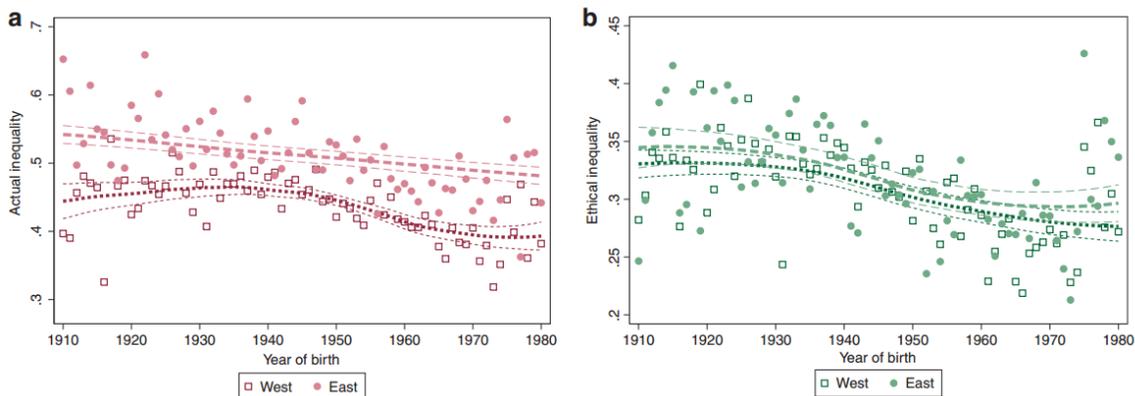


Figure 4: Actual and ethical inequality, as evaluated by Eastern and Western Germans (source: Kuhn 2013)

2.2 Survey experimental designs

Almost all the results in the previous section are built on associations found in cross-sectional studies. It is difficult to draw conclusions about causal relationships from this kind of cross-sectional data. One of the most promising developments in the recent literature is therefore the growing use of experimental designs in large Internet-surveys. In such a design different groups of individuals are subjected to different treatments. The setup is very similar to some of the behavioral experiments that will be discussed in section 4, but in the survey-experimental designs respondents do not always get financial incentives. We briefly discuss two examples.

Cruces et al. (2013) analyze (with a sample of 1060 respondents in Buenos Aires) the importance of the perception of one’s own position in the income distribution. Consistent information on that ranking was provided to a randomly selected treatment group, while perceptions of the control group were not corrected. The results suggest that respondents who overestimate their relative position, are inclined to demand higher levels of redistribution after having been informed that they are in reality poorer than they initially thought. Such experimental study offers more convincing evidence for the importance of the own relative income position than can be obtained with cross-sectional associations, as it can control for confounding factors.

Kuziemko et al. (2015) use a similar method to find an explanation for the puzzle that the desire to redistribute of US citizens apparently has remained stable in an era where inequality has grown considerably. A treatment group receives clear information about the real world developments, a control group does not receive that information. Kuziemko et al. find that giving information indeed changes the general attitudes towards inequality, in the sense that respondents express a stronger desire to reduce inequality, but at the same time the information does not strongly change the degree of acceptance of redistributive government policies.³ In a follow up study, they try to find an

³At least for the income tax. The results for the estate tax are different and further analyzed in great length in

explanation for this fascinating result. Two treatments have a significant effect on the acceptance of redistributive policies: priming respondents to think about the working of the government itself and giving information about specific policy measures. This suggests that the lack of popularity of redistributive policies is driven by the lack of trust in the government (which is more of a problem in the US than in Europe), but that this lack of trust can be partly overcome by drawing attention to the way in which specific policies can indeed be successful in reducing inequality. The political feasibility of redistributive policies clearly does not depend only on the normative attractiveness of prioritarian convictions.

2.3 What did we learn about the acceptance of prioritarianism?

The results summarized in this section are certainly relevant to understand the social reality of redistribution and to gauge the potential of redistributive policies. They offer us a better insight into the acceptance of such policies by the general public and contribute to a better understanding of the underlying differences in the ethical convictions and in the factors explaining these differences. The questions are sufficiently simple to be understood by large fractions of the population and are not too far removed from the cognitions and emotions of the respondents. This is important, as it reduces the danger that the survey itself would induce specific answers from respondents who have never thought deeply about the issue and are not familiar with a more abstract ethical reasoning.

Yet this approach also has its limitations and it certainly is not very informative with respect to the specific questions raised by prioritarianism. The simple questions that are analyzed in these surveys are not explicit about what it exactly means to “reduce inequality”, as they do not specify a measure of inequality. Utilitarians, prioritarians and Rawlsians may all agree that “income differences should be reduced” and the same is true for approaches aiming at equalizing opportunities (in its different interpretations) or for adherents of some sort of capability theory. For a prioritarian who wants to get information about the acceptance of the assumptions underlying this well-defined framework, or about the views of the population on the values of crucial parameters such as the degree of inequality aversion, more detail is needed and the survey questions should be structured in line with the welfare economic theory. Let us now turn to studies that make an attempt in this direction.

3 Questionnaire-experimental work: the acceptance of prioritarianism

The challenge for empirical research on prioritarianism is clear. On the one hand, general questions such as the ones discussed in the previous section, may be close to the intuitions of the respondents, but they are not sufficiently specific to relate the answers to welfare economic theories. On the other

the paper.

hand, welfare economic theory is so abstract that it is nearly impossible to “test” its acceptance among the general population. One cannot ask general questions about the acceptance of prioritarianism versus utilitarianism to subjects who do not have the slightest idea about the meaning of these concepts - and who are likely to have ethical intuitions that are not welfarist. Explaining the relevant concepts before asking about them may induce strong framing effects, as individuals confronted with new information will be influenced by the way this new information is presented. We will first discuss this methodological challenge and then turn to the “testing” of the prioritarian axioms and the “measurement” of the degree of inequality aversion. We will mainly focus on issues related to the income distribution, but also briefly discuss inequalities in health. In subsection 3.4 we summarize some results on the relevance of the veil of ignorance and we end with some results on the limitations of welfarism in general.

3.1 Some methodological preliminaries⁴

Some progress can be made by focusing on the acceptance not of broad theories, but of specific axioms. As illustrated in chapter 2 of this book, the axiomatic approach reduces the intricate reasoning of a complete ethical theory to its essential constitutive building blocks. As Luce and Raiffa (1957, p. 123) write: “By means of a (small) finite number of axioms, we are able to “examine” the infinity of possible schemes, to throw away those which are unfair, and to characterize those which are acceptable.” Focusing on the separate axioms allows the empirical researcher to formulate questions which are theoretically meaningful and at the same time understandable for the population at large. It also allows to focus on well-specified problems, rather than on generalities.

An obvious possibility to test for the empirical acceptance of theoretical ideas or axioms would be to formulate direct questions about them, e.g. “Does an income transfer from a rich to a poor person reduce inequality (or increase social welfare)?” Such direct questions force the respondents to think in an abstract way about general principles concerning problems that are relatively new to them. To avoid framing effects, the largest part of the literature has followed another route. The respondents are then confronted with simple cases or with specific stories that are related to real-world situations, and are asked to give their opinion about what is an optimal (or just, or least unequal) allocation for that specific case. Usually they can choose from a set of precise, quantitatively formulated, options. This circumvents the perception distortions, that have been discussed earlier. Varying the background of the story allows to “test” for context-dependency. Sometimes different variants of the same case are shown to all respondents (*within respondents* design), in other studies different variants are shown to different randomly selected samples of respondents (*between respondents* design). Some of the suggested options correspond to theoretical solutions while others are unrelated to any theoretical concept. The relative acceptance of the various options then offers an indirect, quasi-experimental, way to test different axioms or solution

⁴Methodological issues with this approach are discussed in greater length in Gaertner and Schokkaert (2012)

concepts.

Most (but not all) empirical studies use student samples. Although this choice is mainly explained by the easy availability of students, it has an obvious advantage. Students can be expected to understand complicated questions, which could not be introduced easily in a representative survey. This may be important if one wants to apply the questionnaire approach to discover new theoretical puzzles or to get a better insight into the trade-offs between different axioms. On the other hand, the lack of representativeness of the student samples is clearly problematic if one wants to check whether the theoretical approach is supported by different groups in society, or if one wants to get a better insight into social biases and intercultural differences. There is an obvious trade-off here, but it seems fair to say that future work should try to focus more on representative samples, the more so since the use of Internet has made this practically feasible. There is no reason why one could not set up the type of survey experiments that have been described in section 2.2 in order to test the acceptance of specific theoretical axioms. Quasi-experimental methods can be as easily implemented on the Internet as in a classroom.

We should be aware of the limitations of the simple quasi-experimental setup. It indeed allows us to remain close to the original ethical intuitions of the respondents - but these ethical intuitions may be very different from reasoned opinions. It would therefore be useful to design a dynamic study, in which respondents in a first round are asked simple questions, and in a second round are given additional information about the implications of their choices or about the position that different theories would take. Of course, if we force the respondents to think more explicitly, we will get more sophisticated answers. However, these answers may be influenced heavily by the setup of the questionnaire, and hence by the structure of the dominant theory. This is especially true if the respondents had no strong opinions before they were confronted with the questionnaire. The method to be followed will necessarily depend on what the main objectives of the study are.

3.2 Acceptance of the prioritarian axioms

The seminal work on the acceptance of the different axioms of inequality measurement has been done by Amiel and Cowell and most of this earlier work has been summarized in their 1999-monograph (Amiel and Cowell, 1999). Their original aim was to “test” the popularity of the common axioms of inequality measurement, but in later work they extended their approach to social welfare, poverty, risk, mobility and polarization. We will focus on the results for social welfare, since these are most immediately relevant for this book. Amiel and Cowell investigate various axioms that are relevant for prioritarianism. Anonymity is not tested, but implicitly assumed to hold. We will focus on the Pigou-Dalton transfer principle and on the monotonicity axiom, stating that an increase in the income of an individual leads to an increase in social welfare. As we will see, their results on the transfer principle also throw some light on the separability axiom.

Amiel and Cowell (1999) implement the methodology that we have described above. They focus

on concrete cases, aiming at testing specific axioms. In testing the axioms they predominantly use a within subjects-design.⁵ They go some way in the direction of the dynamic approach that has been suggested, by first asking the respondents to evaluate numerical income vectors, and then confronting them with a verbal formulation of the relevant axioms. The numerical part comes first and after the verbal part, respondents can return and change their answers on the numerical questions if they want to. It turns out, however, that only a small minority of respondents wants to revise their original answers.

Let us illustrate their approach in some detail for the Pigou-Dalton transfer principle. Respondents were shown the following case description:

This questionnaire concerns peoples attitude to income distribution. We would be interested in your views, based on some hypothetical situations. Because it is about attitudes there are no “right” answers. Some of the suggested answers correspond to assumptions commonly made by economists; but these assumptions may not be good ones. Your responses will help to shed some light on this, and we would like to thank you for your participation. The questionnaire is anonymous.

In Alfaland two economic programs are proposed. It is known that both programs will have the same effect on the population except on their incomes and all the people are identical in every respect other than income. In each of the first ten questions there are given two alternative lists of incomes A and B (in Alfaland local currency) which result from these two programs respectively. Please state which program you consider would make the community of Alfaland better off by circling A or B. If you consider that each of the programs is just as good as the other then circle both A and B.

1) A = (1, 4, 7, 10, 13); B = (1, 5, 6, 10, 13).

In the questionnaire, this case (1) was only the first in a list of more than 10 choices. It is a “test” of the transfer axiom. Respondents who accept the latter should prefer society B. Other numerical cases are used to investigate the acceptance of other axioms.

The verbal formulation for the transfer axiom runs as follows:

Suppose there are two economic programs A and B which have only the following difference: The income of person i in program A is x units higher than his income in program B while the income of person j in program A is x units lower than his income in program B. In both programs the income of person i is higher than the income of person j. The incomes of all other people are unaffected by the choice of program A or program B.

- a) Program A would make the community better off.
- b) Program B would make the community better off.
- c) The relative position of other people is also different by A and B. Therefore we cannot say which program would make the community better off.
- d) None of the above.

In the light of the above, would you want to change your answer to question 1?⁶

The approach for the other basic axioms is similar (see Amiel and Cowell, 1999, for more details). Some results are summarized in Table 1. These are based on a sample of 620 students, located in

⁵They opt for a between subjects-design to compare the acceptance of the axioms in different settings, e.g., in measuring inequality versus risk.

⁶“Question 1” was the numerical question given above.

	Numerical	Verbal
Monotonicity	54	55
Transfer principle	47	33

Table 1: Agreement with basic axioms on social welfare (source: Amiel and Cowell, 1999)

Bonn, Koblenz, Stockholm, London and Tel Aviv. The first column gives the agreement with the basic axioms in the numerical formulation of the case, the second column shows agreement with the basic axioms if they are formulated verbally. The degree of rejection of the axioms is striking. This result is not only found in the work of Amiel and Cowell. Similar results are found in a number of other studies (Ballano and Ruiz-Castillo, 1993, Harrison and Seidl, 1994, Leibler et al., 2009, Jancewicz, 2016). The finding that the most basic axioms of prioritarianism are not supported by a large fraction of the population is very robust.

Let us first focus on the transfer principle. In the whole series of studies by Amiel and Cowell, the transfer principle is tested in many different numerical cases. Its rejection is not universal. For instance, a transfer from the highest income to the lowest income in the income vector, is almost generally accepted as welfare improving. The transfer (from A to B) as described in case 1 is one which large groups of respondents (more females than males) do not see as a welfare improvement. Note that this is not necessarily irrational. The psychological mechanism at work is nicely illustrated in Figure 5, which is also taken from Amiel and Cowell (1999). The transfer in case 1 lowers the distance between the second and the third individual in the row, but it increases the distance between the worst-off and the second worst-off. Respondents that focus on that distance may (and apparently do) see the transfer as welfare-decreasing. These respondents appear (a) to focus on the situation of the worst-off, and (b) to have the intuition that the welfare of the worst-off decreases if the income distance with the second worst off increases, despite the fact that his/her own income remains the same. It is interesting to notice that these results suggest that respondents do not consider the pair of individuals that are involved in the transfer in isolation, but evaluate the transfer taking into account the income of other individuals. This amounts to a rejection of the separability axiom that is essential for prioritarianism.

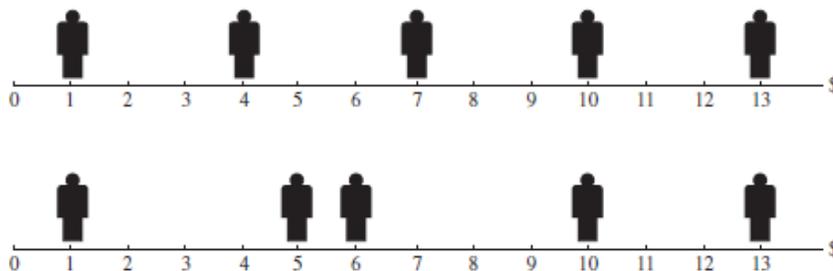


Figure 5: Transfer principle (source: Amiel and Cowell, 1999)

Let us now turn to the interpretation of the rejection of monotonicity. Respondents are asked to evaluate an income increase for one individual and the questionnaire does not make explicit the relationship between income and individual well-being. Monotonicity is mainly rejected when an income in the upper part of the income vector increases. A natural interpretation of this result is again that respondents believe that a deterioration of the relative position of the poor leads to a decrease in their well-being. That the relative income influences individual well-being is also found in other questionnaire studies (see, e.g., Konow, 2001, Johansson-Stenman et al., 2002).⁷ Rejecting monotonicity in the space of incomes therefore does not necessarily imply that the Pareto principle in the space of individual well-being is also rejected (see Amiel and Cowell, 1994). On the other hand, other questionnaire studies that have formulated the question explicitly in utility terms also have found limited support for the Pareto-principle (Mc Clelland and Rohrbaugh, 1978).

This discussion on the interpretation of the axioms points to a more general issue. Although focused on social welfare evaluation, the questions in the work of Amiel and Cowell and in related work are about income redistribution, not about welfare redistribution. Hence, this work does not allow to compare the acceptance of, e.g., utilitarianism and prioritarianism, since redistributive transfers will also be desirable for a utilitarian (if individual well-being is concave in income). We will come back to this point in section 5.

3.3 The degree of inequality aversion

Some authors have (implicitly or explicitly) taken the utilitarian or prioritarian framework as given and have focused on estimating the degree of inequality aversion within that framework. Contrary to the prioritarian approach advocated in this book, this literature does not aim at measuring inequality aversion in the space of individual levels of well-being, but focuses on specific components of well-being. We first summarize some findings on the income distribution and then raise the issue of health inequalities.

3.3.1 Income inequality

An early example of this work is by Glejser et al. (1977) who let students choose between different allocations of scholarships. Broadly speaking, in the more recent literature this question has been tackled in two ways: either through the idea of a leaky bucket (e.g. Amiel et al., 1999) or through a straightforward comparison of two distributions (e.g. Carlsson et al., 2005). Pirttilä and Uusitalo (2010) compare the results obtained with the two methods for a representative survey of the Finnish population. We present their formulation of the questions in some detail. The other studies follow a similar approach, although the specific formulation and the underlying formal model are different.

⁷We are referring here to studies that explicitly focus on the *normative* evaluation of income distributions. Of course, the happiness literature contains an overwhelming amount of *positive* evidence that the relative position of individuals in the distribution of income (or well-being) matters for their own subjective well-being (see, e.g., the survey by Clark and D'Ambrosio, 2015).

Pirttilä and Uusitalo (2010) assume that social welfare can be represented by an Atkinson-type social welfare function⁸

$$S = \sum_i \frac{y_i^{1-e}}{1-e} \text{ for } e \neq 1, \quad (1)$$

$$S = \sum_i \ln y_i \text{ for } e = 1. \quad (2)$$

They implement the leaky bucket-method by asking each respondent one question, corresponding to one value of e . For $e = 1$, the question goes as follows:

What is your opinion of the following reform proposal?

The taxation of all high-income earners, whose disposable income exceeds €3300 per month, is increased. The money is spent for the benefit of those low-income earners whose disposable income is less than €800 per month. The high-income earners can, however, react to the tax increase by reducing their work effort, and part of the money goes to administrative expenses. Therefore, for each 100 paid by the high-income earners, only 25 can be spent for the benefit of low-income earners.

Are you, nevertheless, in favor of this proposal?

1. Yes 2. No 3. Cannot say.

The incomes of €3300 and €800 correspond respectively to the average income in the highest and the lowest decile of the Finnish income distribution of 2003, inflated to 2006 values. A respondent with a personal value of $e > 1$ should be in favor of the proposal, someone with $e < 1$ should reject it. Similar questions, corresponding to $e = 0.5, 2$ and 3 are asked to other respondents. Pooling all the answers makes it possible to estimate the distribution of e in the population. Since a majority of respondents rejects the proposal in the case $e = 0.5$, it appears that the median inequality aversion of the respondents lies below 0.5 .⁹

In the second approach, the respondents get a choice between two distributions. Again, each respondent is asked one question, corresponding with the same value of e as was implemented in his/her leaky bucket question. The income distributions are again calibrated on the real Finnish situation and for $e = 1$ the question then becomes:

Let us imagine that in wage negotiations two different alternatives are considered. Which of the following do you prefer?

1. If all employees are ordered from the lowest-income earner to the highest-income earner, someone belonging to the lowest decile earns 1570 in a month, a person with average income earns 2340 and a person belonging to the highest decile earns 3480.

2. Income differences rise and the average income is increased, so that the low income earner gets 1280 per month, the person with average income gets 2580 and the high-income earner 5190.

3. Cannot say.

⁸To avoid confusion with the other chapters, we use the symbol e to measure the degree of aversion to income inequality, rather than γ that is used to denote the degree of aversion to inequality in well-being.

⁹The fraction of the sample rejecting the proposal for values of e larger than 0.5 is even larger, as it should be if the answers are consistent with the underlying theory.

A respondent with $e = 1$ should be indifferent between the two distributions, someone with $e > 1$ should prefer distribution 1 and someone with $e < 1$ should prefer distribution 2, since for the latter the increase in average income is large enough to compensate for the increased inequality. In this case, the more equal distribution is preferred by all respondents, suggesting that the median inequality aversion is larger than 3.

The difference between the results obtained with the two methods is surprisingly large, and this is even more striking because similar differences are found in other studies. Different explanations can be proposed for this difference. First, the reference points in the two questions of Pirttilä and Uusitalo are different. The leaky bucket question describes a tax increase from an actual situation that is more unequal, the formulation of the comparison question suggests that the less unequal distribution is the starting point. The different answers could then just reflect status quo-bias. This raises the question whether this status quo bias is a genuine psychological phenomenon that also would manifest itself in real world decisions or just a framing effect, due to the questionnaire setting. If the former holds true, it is still another question whether this status quo bias has normative significance, e.g. in a libertarian setting with respect for property rights, or should just be seen as a normatively irrelevant “true bias”.

While the different reference points may contribute to the explanation for the different results obtained with the two methods in the setting of Pirttilä and Uusitalo, this is not sufficient to explain why a similar difference is found in the other studies with different formulations. Pirttilä and Uusitalo themselves propose an alternative explanation. They suggest that the degree of inequality aversion may be context-dependent - and different in a tax setting drawing attention to the cost induced by reducing work effort compared to a setting with two apparently pre-tax income distributions. Intuitions in the former case may be driven by the conviction that effort should be rewarded, while the different distributions in the latter story may be interpreted as a matter of bargaining. In fact, this explanation is closely related to the hypothesis about the lack of trust in the government that we have described in section 2.2. Of course, all these explanations suggest that the ethical intuitions of respondents cannot be fully captured within an outcome-related approach. We will return to this point in section 3.5.

3.3.2 Health inequality

Until now we focused on the inequality aversion with respect to incomes. There is also a literature on inequality aversion in the space of health. This inequality aversion is crucial in a prioritarian approach to health care priority setting. The dominant approach in cost-effectiveness studies on health implements an objective function which is an unweighted sum of so-called “qaly”s, quality-adjusted life years. It is not always clear whether qaly’s should be seen as a measure of health or as a measure of utility (see Hausman, 2015, for a critical discussion), but this issue goes beyond the topic of this chapter. What is relevant for our purpose, however, is the near consensus among

the respondents in survey studies that this generally applied “unweighted” (quasi-utilitarian) objective function is not acceptable and that some (prioritarian) distributional weighting should be introduced (see Gaertner and Schokkaert, 2012, and Costa-Font and Cowell, 2019, for overviews of this literature).

A limited number of studies (e.g. Edlin et al., 2012; Ottersen et al., 2014) present measures of the inequality aversion in this health setting. These studies (with a design similar to the comparison method described in the previous subsection) find values of health inequality aversion that are in the same range as those of the income inequality aversion found with the same method, i.e. larger than 1 but smaller than 10.

Yet the majority of the studies do not focus on pure health inequality but rather on income (or socioeconomic status)-related health inequality (e.g. Dolan and Tsuchiya, 2011, Robson et al., 2017, Cookson et al., 2018, Lal et al., 2019). In this setting, Robson et al. write the social welfare function (for $\epsilon \neq 1$) as¹⁰

$$S = \frac{1}{1 - \epsilon} \sum_k f(x_k) H_k^{1-\epsilon} \quad (3)$$

where H_k is the health level of income group k , and $f(x_k)$ the proportion of the population in that group. In this framework the estimated degree of inequality aversion varies but is usually much larger than in the “pure” health inequality case. Dolan and Tsuchiya report a value of 29, and in other studies the degree of inequality aversion cannot even be estimated because preferences are no longer monotonic (Abasolo and Tsuchiya, 2004, 2013). This income related inequality aversion is a different concept than the pure health inequality aversion, however. The way the questions are formulated draws attention to the issue of cumulative deprivation: respondents have to compare (changes in the) distributions of health in which the worse off are not only less healthy, but also have a lower income. It is not really surprising that they then show a more outspoken preference for redistribution.¹¹

Robson et al. (2017) and Cookson et al. (2018) offer an interesting example of the dynamic approach that we advocated above. In their study (with admittedly a small number of respondents) they design two feedback procedures: one in which respondents receive more information on the implications of their choices in the first round, and a second one in which respondents are shown a video in which different characters introduce relevant arguments concerning the trade-off between total health and its distribution. Their hypothesis is that the simple questionnaire setting induces a pro-egalitarian attitude among respondents and that they are likely to become less inequality averse if they receive more information. This hypothesis is corroborated in their study in which the income related health inequality aversion is around 10. Of course, compared to the traditional

¹⁰We focus on the Atkinson type of social welfare function, since this yields results that are comparable between the different studies and domains. Robson et al. (2017) also show results for a Kolm-type social welfare function.

¹¹The issue of the correlation between income and health has also been tackled in the context of inequality measurement, with special focus on the relevance of the concentration index - see, e.g., Bleichrodt et al. (2012) and Tarrow (2015).

values in a pure income (or health) setting, this is still very large.

3.4 The veil of ignorance, ex ante and ex post

At least since Harsanyi (1953, 1955) and Rawls (1958, 1971) the concept of the “veil of ignorance” (VOI) has played a prominent role in the philosophical literature on distributive justice. It is built on the assumption that individuals would be free of any personal bias when they are put in the hypothetical situation in which they do not know their own relative position in society. Choices behind the veil would therefore yield an acceptable theory of justice and of social welfare. Different philosophers, notably Harsanyi and Rawls, make different assumptions about the “thickness” of the veil, i.e. the nature and the amount of information that is hidden, and derive different conclusions from it. For Rawls, decision makers under “his” VOI would choose the maximin criterion, for Harsanyi a rational decision maker behind the VOI would necessarily opt for utilitarianism, and not for prioritarianism. It is therefore relevant to check if the VOI is accepted by respondents as an intellectual device to arrive at a theory of social welfare. In fact, as we will see in section 4, the VOI has often been used in laboratory experiments precisely for that purpose. It is less common in the questionnaire literature however.

The most relevant study on the topic has been conducted by Amiel et al. (2009). They reformulate the Alfaland-type question described above and introduce the idea that there are five regions in Alfaland, each with a different income. Respondents are again asked to indicate which of two situations (income vectors) would result in a better situation for Alfaland. There are two versions of the questionnaire. In the first, respondents are asked to put themselves in the shoes of an outside observer. In the second, they are asked to imagine that they themselves are assigned to one of the regions of Alfaland, with equal probability of finding themselves in one of the five regions. This second option is of course the VOI-treatment. The results are mixed (because students in Israel do not follow the general pattern), but in broad lines the authors find that the standard prioritarian axioms, and more specifically the Pigou-Dalton transfer principle, are followed to a larger extent in the “ethical observer” treatment than in the VOI-treatment.

Bosmans and Schokkaert (2004) find similar results in a different setting. Students are asked to evaluate the situation of two firms that each plan to employ 100 recently graduated students. In each firm there are three types of jobs, corresponding with different incomes. Bosmans and Schokkaert have a number of different treatments but two of them correspond to those proposed by Amiel et al. (2009): the “impartial and sympathetic observer” and the “rational individual behind the veil of ignorance“. Only between 10 and 13% of the respondents respect the axioms of expected utility theory, which in the VOI-treatment is the basic justification of utilitarianism by Harsanyi. Moreover, as in Amiel et al. (2009), the response patterns are significantly different between the observer- and the VOI-treatments. Similar results are reported by Bernasconi (2002).

While the previous studies were set in a setting with incomes, Pinto-Prades and Abellán-

Perpiñán (2005) implement the VOI to derive the social values attached to different health states, and compare the obtained results with those from a so-called “person trade-off” approach (Nord, 1995), which focuses on the trade-off between the numbers of individuals treated and the severity of their disease. This is a typical “ex post” approach, in which the respondent is put in the position of a decision-maker who has to take allocation decisions in front of the veil of ignorance. The survey was administered to a stratified random sample of 300 respondents, drawn from the general population of the Barcelona area. The two treatments again lead to very different results. The “outside observer” in the person trade-off case finds it more important to avoid that people die. Again, the large differences between the person trade-off and the VOI confirm that it is not acceptable to suppose that decisions behind a veil of ignorance would coincide with the decisions taken by an ethically motivated outside observer. This important finding seems to be quite robust in a large variety of settings.

3.5 A broader set of questions: welfarism, desert and taxation

The results described until now suggest that the prioritarian framework is not generally accepted by the population. A strong indication of this is the large degree of rejection of the transfer principle. In this section we collect some challenges raised by the questionnaire studies that are even more basic, because they put in doubt the traditional welfarist assumptions underlying prioritarianism.

First, there is by now overwhelming evidence that for (almost all) respondents the causes of income differences play an essential role in the evaluation of their legitimacy. As a clear example, there is near consensus that effort should be rewarded, and that income differences due to brute luck should be compensated. We mentioned already that this is a quasi-universal finding in the large survey studies. It is also found in the questionnaire-experimental work.¹² Surveys of these results can be found in Konow (2003) and Gaertner and Schokkaert (2012, chapter 4). In fact, already in the seminal study in this line of research, Yaari and Bar-Hillel (1984) found that Israeli students made a distinction between welfare differences due to needs and to tastes, and considered redistribution to be more legitimate in the former than in the latter case. All these results can be seen as a clear rejection of welfarism and as supportive evidence for the popularity of the notion of equality of opportunity. The relationship between this notion and prioritarianism should be further explored. One interesting avenue is suggested by the theory of fair allocation, that integrates ideas about responsibility and compensation in its normative approach to individual well-being (see, e.g., Fleurbaey and Maniquet, 2011, 2018).

Second, when thinking about taxation, many (or even most) respondents go beyond the simple ideas underlying the optimization of a traditional social welfare function. This is convincingly shown by the work of Weinzierl with Internet samples of US respondents. Weinzierl (2004) asks the respondents to rank different tax policies (see Figure 6). The equal sacrifice policy (in his case

¹²And section 4.5 will illustrate that it is also an important element in behavioral experiments.

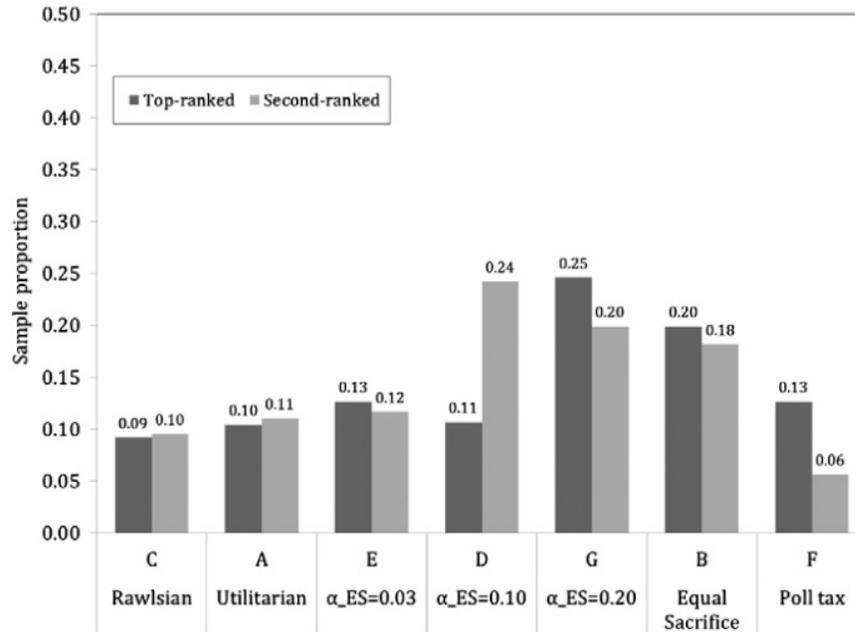


Figure 6: Utilitarian versus equal sacrifice taxation (source: Weinzierl, 2004)

boiling down to proportional taxation) is more popular than either the Rawlsian or the utilitarian policy. Even more popular is a mixed approach with α_{ES} (the weight given to equal sacrifice) equal to 0.2. In another paper, Weinzierl (2017a) shows that respondents resist the fully egalitarian solution, even if there no incentive problems and the initial income difference is due to brute luck.¹³ Moreover, a large share of respondents prefer a classical benefit-based logic over the prioritarian logic of diminishing marginal social welfare. These two views are linked: respondents who more strongly resist equalization are more likely to prefer the classical benefit-based principle. As a last example, Weinzierl (2017b) asks respondents to choose between situations with explicit information about monetary income and well-being (see Figure 7). He suggests that the loss in well-being of person F in situation 2 is due to envy. About half of his respondents does not accept this as a legitimate reason to compensate person F. This study is one of the few examples where well-being is explicitly introduced in the formulation of the questions. Note that the finding that the negative effect of growing inequality on the well-being of the worse-off should not be compensated, clashes with some of the findings that we discussed in section 3.2. All the results of Weinzierl are obtained with US-respondents and surely when questions are asked about specific policies, such as income taxation, respondents may be strongly influenced by the institutions in which they are living. This is illustrated by comparing Weinzierl’s results to those of Tarroux (2019). In his French sample, respondents are first asked to rank final income distributions. When they are in a second stage

¹³This case is analyzed theoretically in Saez and Stantcheva (2013), and they show survey results that go in the same direction as those of Weinzierl (2017a).

Now we will ask for your opinion regarding this scenario.

The table below is the same as what you saw on the previous screens. Please think about the two situations from the perspective of an objective observer, for example as if you were a policymaker.

Given the information in the table, which situation do you think is better?

	Situation 1	Situation 2
Money		
Person F:	5,000	6,000
Person G:	10,000	16,000
Total	15,000	22,000
Well-being		
Person F:	120	110
Person G:	160	166
Total	280	276

Please choose the statement below that best describes your opinion.

Strongly prefer Situation 1 Slightly prefer Situation 1 Slightly prefer Situation 2 Strongly prefer Situation 2

Figure 7: Choice between income/well-being bundles (source: Weinzierl, 2017b)

informed about the average tax rates leading to these final distributions, the majority changed their mind and opted for a progressive taxation policy. This means that not only the final outcomes matter, but also the shape of the tax function itself. Again, this finding cannot immediately be taken up in the simple prioritarian framework.

4 Ethical preferences in the lab

We now turn to an overview of the findings obtained from lab experiments. We will cover the same topics as in the previous section. After some methodological considerations (section 4.1), we first discuss the acceptance of the basic prioritarian axioms (section 4.2) and the estimation of the degree of inequality aversion (section 4.3). We then show the results of recent work on the choice between ex ante and ex post perspectives in a context of risk and uncertainty (section 4.4) and finally turn to some broader questions concerning the popularity of welfarism (section 4.5).

4.1 Methodological considerations

While there is a huge experimental economic literature on social preferences (altruism, reciprocity, envy, etc.)¹⁴, ethical preferences have received less attention. Two methods, that we already have introduced in section 3.4 on questionnaire studies, have been investigated in the literature.¹⁵ The first method puts a subject in a position to choose a payoff allocation for *other* subjects in the lab.

¹⁴See for instance Andreoni and Miller (2002), Charness and Rabin (2002), Engelmann and Strobel (2004), Fehr et al. (2006) and Cooper and Kagel (2015) for a survey.

¹⁵See Aguiar et al. (2013) for a comparison between different methods.

That is, subjects are grouped by $(n + 1)$ and one of them, called the *social planner*¹⁶ is asked to allocate a certain amount of money among the n others, called *stakeholders*. That *real* stakeholders are affected by the choice made by the social planner should make her aware of her responsibilities and works as an incentive mechanism.¹⁷ The subjects playing the role of social planner, receive a payoff that is not affected by their choices. It can then be assumed that her choices reveal her view about how a resource should be shared among people.¹⁸

According to the second method, the subject is involved in the situation, i.e., she is a member of the group for which she is asked to choose an allocation of payoffs, but her decision is made *under a veil of ignorance*. Almost all the experiments on which we will report, use a design in which subjects know the probability of each position. For instance, the subject is grouped with $(n - 1)$ other subjects and she is asked to choose between (x_1, \dots, x_n) and (y_1, \dots, y_n) knowing that the probability of receiving x_i is equal to $\frac{1}{n}$. An alternative is to impose uncertainty, that is, subjects know the possible levels of payoff that they can receive but do not know the probability of each of them. Even if one accepts such a device as a way of determining what is fair or not, an open question is to know if such a framework works in the lab. As we have seen, it is not certain that a veiled choice of a group allocation really reflects ethical preferences.

Some papers in political science use a variant of the veil of ignorance approach in which groups deliberate about the principles of justice and next make a decision (Frohlich et al., 1987a, 1987b, Frohlich and Oppenheimer, 1990). Their aim is to test the Rawlsian “original position” and its implications: do people under the condition of a veil of ignorance converge to a consensus about what is fair, and is this consensus view to maximize the well-being of the worst-off?

4.2 Acceptance of the prioritarian axioms

Only few experimental papers aim at testing prioritarian axioms, even though there are a number of studies that investigate people’s views of distributive justice. It seems there is no clear consensus in the literature. Some papers, like Frohlich et al. (1987a, 1987b), Herne and Suojanen (2004) or Traub et al. (2005), find a rejection of maximin or leximin views and a support for maximizing the average income with a floor income guarantee. For instance, in Frohlich et al. (1987b), subjects are invited to discuss four alternative views: maximin, maximizing the average income, maximizing the average income while guaranteeing everybody a floor income and maximizing the average income with a range constraint. Their main results are that all the groups reach a consensus and that the most preferred option is to maximize the average income with a floor income (25 over 27 groups, 109 over 149 individuals). Yet other studies have found that respondents express a strong desire to redistribute income (e.g., Dickinson and Tiefenthaler, 2002; Becker and Miller, 2009; Muller and

¹⁶Subjects are not told that they are a social planner or stakeholder. Experimentalists are used to adopt a neutral language in order to avoid any demand effect.

¹⁷In Traub et al. (2005) and Traub et al. (2009), the social planner is pushed to the fore in order to increase her sense of responsibility.

¹⁸This assumption seems to be empirically valid: see for instance Capellen et al. (2013).

Renes, 2017). Based on a large and heterogeneous pool of German respondents, Muller and Renes (2017) found that the egalitarian allocation is the most popular allocation (50% of subjects opt for this allocation) but the maximin solution also remains attractive (40% of subjects). Only a low proportion of subjects (10%) seek to maximize total income.

Even if such studies are informative, they do not test the prioritarian axioms. To the best of our knowledge, the only lab-experiments that are relevant for this purpose are reported by Beckman et al. (2002) and Traub et al. (2009). The former propose a test of the Pareto axiom while the goal of the latter is to study the popular support for the axioms of Pareto and Pigou-Dalton transfer and for the Lorenz and Generalized Lorenz dominance criteria. None of them are trying to test the axioms of anonymity, separability and continuity.

Beckman et al. (2002) test the acceptability of the Pareto axiom according to which social welfare increases if the income of at least one person increases while that of others remains at least unchanged.¹⁹ In their experiment, subjects are grouped by five and are asked to vote between two payoff allocations (one obtained from the other by an increment of payoff of one or several individuals) without and with knowledge of their position in the group. Using a pool of subjects in different countries (USA, Russia, Taiwan and China), the authors found that the Pareto axiom is well supported by subjects under the veil of ignorance: on average, the Pareto axiom is rejected in only 10% of the relevant comparisons. This rejection rate varies from 5.4% in Taiwan to 20% in Russia. In the treatment in which positions are known, the rejection is higher: from 6.8% in the USA to 28% in China.

Traub et al. (2009) report the results of an experiment which aims at testing the usual dominance criteria for comparing income distributions. More precisely, they investigate the support for the axioms of Pareto, Pigou-Dalton, Lorenz and Generalized Lorenz dominance in the domain of income. The Lorenz criterion states that an income vector \mathbf{v} is preferred to a distribution \mathbf{w} if the share of total income belonging to the $p\%$ poorest individuals in \mathbf{v} is at least equal to that in \mathbf{w} for any $p \in [1, 100]$, and strictly larger for at least one $p \in [1, 100]$. The Generalized Lorenz dominance criterion works the same except that the share of total income is replaced by the amount of income. That is, \mathbf{v} is preferred to \mathbf{w} if the total income belonging to the $p\%$ poorest individuals in \mathbf{v} is at least equal to that in \mathbf{w} for any $p \in [1, 100]$, and strictly larger for at least one $p \in [1, 100]$. It can be shown that: (i) acceptance of Pareto and Pigou-Dalton implies acceptance of Generalized Lorenz dominance;²⁰ (ii) acceptance of Pigou-Dalton implies acceptance of Lorenz dominance.

Subjects are asked to give a complete and strict preference ordering of 12 income distributions. The authors run four treatments varying in terms of the subject's position. We choose here to focus

¹⁹This axiom was called "monotonicity" in the previous section. We follow here the Beckman et al.-terminology.

²⁰More precisely, for two *income* vectors with the same number of individuals, the following conditions are *equivalent*: (i) \mathbf{v} dominates \mathbf{w} according to the Generalized Lorenz dominance; (ii) \mathbf{v} can be obtained from \mathbf{w} by a series of Pigou-Dalton transfers of income and/or increments of income (Pareto) and/or permutation of individuals; (iii) $\sum_i \phi(v_i) \geq \sum_i \phi(w_i)$ for all ϕ such that $\phi' \geq 0$ and $\phi'' \leq 0$ and $\sum_i \phi(v_i) > \sum_i \phi(w_i)$ for at least one ϕ in this class of functions.

on the two of them that are most relevant to investigate ethical preferences: when the subjects rank distributions under a veil of ignorance (*individual-choice* treatment) and when they play the role of a social planner (*social planner* treatment).²¹

No	Vector of incomes	Individual-choice		Social planner	
		Mean	Median	Mean	Median
1	(30, 30, 30, 30, 30)	5.5	6	5.3	6
2	(25, 27.5, 30, 32.5, 35)	6.38	7	5.88	6
3	(20, 25, 30, 35, 40)	6.45	6	6.03	6
4	(20, 20, 30, 40, 40)	5.32	5.5	5.48	5
5	(20, 30, 30, 30, 40)	5.02	5	5.89	7
6	(5, 10, 30, 50, 55)	3.52	2	3.38	2
7	(5, 30, 30, 30, 55)	3.03	2.5	3.62	2.5
8	(35, 35, 50, 55, 60)	9.40	10	8.65	10
9	(35, 35, 35, 45, 90)	9.67	10	8.8	10
10	(7.5, 7.5, 50, 55, 60)	4.05	3.5	4.06	3
11	(7.5, 7.5, 35, 45, 90)	4.28	4	4.88	4.5
12	(0, 30, 40, 125, 125)	3.37	1	4.02	2

This table reports the payoff distributions and the mean and median Borda counts for individual-choice and social-planner treatments. Note that the subjects are presented with distributions of $\times 1000$ income in order to increase the .

Table 2: Distributions and their Borda score (Traub et al., 2009)

Dominance criteria	Nb of cases	Acceptance rate (%)	
		individual-choice	social planner
Pareto-rank dominance	15	91.8	82.5
Progressive transfer dominance	17	57.8	61.6
Lorenz dominance	53	64.4	61.9
trade-off between efficiency and equity	25	59.1	55.2
no trade-off	28	75	67.9
Generalized Lorenz dominance	41	75.1	71.8
trade-off (no Lorenz dominance)	13	88.2	80.1

The acceptance rate is defined as the mean proportion of responses corresponding to a given dominance criterion.

Table 3: Acceptation rates of basic axioms (Traub et al., 2009)

The second column of Table 2 displays the distributions the subjects are asked to rank. Distri-

²¹In the two other treatments, subjects are asked to rank lotteries without any reference to an income distribution or are put in the position of a group dictator evaluating distributions in which she is involved. The incentive mechanism in the experiment works as follows. First, in the individual-choice treatment, five subjects in the room are randomly selected for payment; in the social-planner treatment, 6 subjects are randomly selected: one for being the social planner (who is visible to all other subjects) and five for payment. Second, two distributions are randomly chosen from the set of the 12 available distributions and the preference ranking over these two distributions is applied for payment.

butions 1 to 7 have the same mean income and distributions 2-7 can be obtained from distribution 1 by a Pigou-Dalton transfer of incomes. Distributions 8-12 have a higher mean income than distributions 1-7 but a higher level of (relative) inequality. Note that distributions 8 and 9 are the minimal income maximizing distributions. The second column of Table 3 reports the number of dominance relationships for each criterion. Interestingly, the authors study how people arbitrate between efficiency and equity. There are 25 cases in which a distribution dominates another one according to the Lorenz criterion but not according to the Generalized Lorenz one (e.g., distribution 8 Lorenz-dominates distribution 9 whereas the Generalized Lorenz criterion is silent). Among the Generalized Lorenz dominance relationships, 13 of them are in conflict with inequality, e.g. inequality is higher in distribution 8 than in distribution 2 according to relative Lorenz dominance but the former dominates the latter according to generalized Lorenz dominance.

The experiments were conducted at the University of Hannover with 252 students (between 60 and 66 subjects per treatment). Table 2 provides the mean and median Borda counts of each distribution and Table 3 summarizes the support of the basic dominance criteria in the two treatments. First, it is found that the Pareto principle is largely accepted by the subjects: the subject's choices respect Pareto dominance in 80% of the 15 comparisons.²² This confirms the finding of the Beckman et al. (2002) study, but conflicts with the dominant findings in the questionnaire studies (see section 3.2).

Second, and better in line with the questionnaire studies, the Pigou-Dalton axiom receives relatively low support. The acceptance rate is about 60% in both treatments. Considering distributions 1 to 7 (with the same mean), the aggregate ranking based on (mean) Borda counts is at odds with the ranking according to Pigou-Dalton axiom: $3 \succ 2 \succ 1 \succ 4 \succ 5 \succ 6 \succ 7$ for the individual choice treatment and $3 \succ 5 \succ 2 \succ 4 \succ 1 \succ 7 \succ 6$ for the social planner treatment against $1 \succ 2 \succ 5 \succ 3 \succ 4 \succ 7 \succ 6$ according to the Pigou-Dalton criterion.

Third, Lorenz dominance also receives low support but is more accepted when there is no conflict between efficiency and inequality. At the same time, Generalized Lorenz dominance is well supported.

The last observation pushes the authors to suggest that subjects have a strong preference for efficiency. However, one could as well argue that subjects (also) put a greater weight on the worst-off members. Indeed, the mean and median Borda counts indicate that the minimal income maximizing distributions (distributions 8 and 9) are the most preferred. Moreover, when comparing distributions 10-11 with distributions 1-5, the latter seem to be preferred to the former although the former have higher average earnings and the Generalized Lorenz dominance criterion is silent.

All in all, these two papers tend to reject one the main axiom of prioritarianism, namely the Pigou-Dalton axiom applied to income distributions. This is in line with the questionnaire findings. More surprisingly, the Pareto principle seems to be well accepted in the laboratory.

²²Since this proportion is similar irrespective of the potential conflict with inequity, the authors argue that the non-respect of the principle is due to errors.

4.3 The degree of inequality aversion

Some papers studied more deeply how people arbitrate between equity (viewed as maximin) and efficiency (viewed as total income). Schildberg-Hörisch (2010), Frignani and Ponti (2012) and Heufer et al. (2018) try to disentangle the role of ethical and risk preferences behind a veil of ignorance and provide an estimate of inequality aversion. The general framework is a dictator game with uncertainty, i.e., subjects are asked to make a choice of a pair of payoffs within various budget sets:²³

$$z = qy_1 + y_2 \quad \text{with } q \geq 1 \text{ and } z > 0 \quad (4)$$

where y_1 and y_2 are the low and high payoffs, z the total expenditure and q the cost of averaging payoffs. Three types of choice problem can be introduced. In a *risk problem*, the pair (y_1, y_2) is defined as a pair of payoffs in a bad and high state, the probability of each state being 50%. In a *choice under a veil of ignorance*, (y_1, y_2) is the sharing of the total payoff between two persons given that positions have the same probability. Schildberg-Hörisch (2010) and Frignani and Ponti (2012) also run a *standard dictator* treatment where the subjects know their position (as introduced by Andreoni and Miller, 2002).

In Schildberg-Hörisch (2010) the dictator has an endowment of 12 experimental units and has the opportunity to transfer money to the receiver knowing that the transfer is costly (1 unit transferred has a cost of 0.5). Using a pool of 167 subjects from the University of Mannheim, the author found that only 14% of subjects exhibit maximin preferences. On average, subjects send 3.5 points from the dictator to the receiver (i.e., 30% of the total payoff) entailing a decrease in total payoff from 12 to 10.25 points. Interestingly, she also reported that choices under the veil of ignorance are not (only) explained by risk aversion since subjects send 3.1 points from the good to the bad state in the risk problem. While 44% of the subjects transfer the same amount in the two types of decision problem, 35% of subjects transfer more in the allocation choice under the veil of ignorance. Finally, in the standard dictator game, subjects transfer 1.6 points, significantly less than when positions are uncertain.

In the experiment conducted by Frignani and Ponti (2012) subjects face 24 decision problems in one of the three conditions described above. Based on a pool of 192 students of Universidad de Alicante, they estimate a simple mean-variance utility function:

$$u_i = \bar{y} - \delta\sigma(y) \quad (5)$$

The estimated values of δ are: 0.267 in the standard dictator game, 0.347 under the veil of ignorance,²⁴ and 0.316 in the risk problem. Importantly, they found that there is no statistically

²³This experimental design was introduced by Andreoni and Miller (2002) and Fisman et al. (2007) to study individual social preferences.

²⁴For the sake of illustration, with this degree of inequality aversion, (70, 32) is preferred to (44, 44).

significant difference between the values of δ based on allocation decisions under the veil of ignorance and in the risk problem. That suggests that, on average, decisions made under the veil of ignorance are mainly explained by risk aversion. However, the average values may hide individual heterogeneity, as was found by Schildberg-Hörisch (2010).

Heufer et al. (2018) report an experiment where subjects face 40 decision risk problems and 40 allocation problems under a veil of ignorance. The use of a within-subject design allows them to take choices in the risk problems as a benchmark to evaluate ethical trade-offs. For instance, if the chosen payoff of the poor is higher (lower) than the one of the bad state, the decision-maker reveals a preference for equality (efficiency). The experiment was conducted in Xiamen University with 92 subjects. We present here some results that may be of interest for the evaluation of prioritarianism.

First, the authors found that the mean transfer to the poor/bad state amounts to 27-28% of the total payoff in both types of decision problem. However there is strong heterogeneity between subjects: for 51 of the 92 subjects, the authors cannot reject that they sent the same amount to the low payoff in both treatments. These subjects are classified as socially agnostic. Twenty-four subjects can be classified as equity preferring in the sense that the amount sent to the poor is higher than the one sent to the bad state. The other seventeen subjects are efficiency-preferring.

Second, the authors estimate the demand functions for the low payoff (see Table 4). A rather surprising result is that in the VOI-treatment the demand of the equity-preferring subjects is more sensitive to price than that of the efficiency-preferring subjects. But, as the authors point out, an increase in the expenditure level z (i.e. the total amount to share between qy_1 and y_2) increases the low payoff by a higher amount for the equity-preferring subjects than for the others. Following an increase in z by 10 units, an equity-preferring subject increases the low payoff by 2.4 units, while a subject focusing on efficiency increases it by 1.4 units.

	Types of subjects		
	Equity	Agnostic	Efficiency
Constant	28.32**	29.79**	28.45**
Price q	-9.57**	-10.06**	-10.24**
Expenditure z	0.20**	0.20**	0.21**
Constant \times VOI	5.13**	0.08	-9.06**
Price $q \times$ VOI	-1.93**	0.11	3.71**
Expenditure $z \times$ VOI	0.04**	-0.004	-0.07**
Obs.	1,920	4,080	1,360
R ²	0.72	0.62	0.58

Dependent variable: payoff of the poor/bad state. OLS regressions with robust standard errors clustered at the individual level.

Table 4: Demand for the low payoff (Table 2, Heufer et al., 2018)

Third, the authors estimate a structural model assuming that a decision-maker maximizes the

following subjective utility function:²⁵

$$V(y_1, y_2) = \alpha \frac{(y_1)^{1-\gamma}}{1-\gamma} + \frac{(y_2)^{1-\gamma}}{1-\gamma} \quad (6)$$

where γ reflects subject's inequality or risk-aversion and α is the ratio between the subjective probability of receiving the low payoff and that of receiving the high payoff ($\alpha > 0$). A value of α larger than 1 means that the probability of the low payoff is larger than that of the high payoff.²⁶ The case where $\alpha = 1$ corresponds to the expected utility criterion. In a risk problem, γ may be interpreted as the degree of risk aversion. In an allocation decision under a veil of ignorance, it may capture risk and inequality aversion. The authors estimate equation (6) for each subject. Table 5 reports the mean and median values of γ and α . The level of γ is higher in the allocation problem under a veil of ignorance than in the risk problem. Again they found a quite large heterogeneity between subjects, in particular under the veil of ignorance. The comparison of individual estimates of γ in the two contexts allows the authors to identify the importance of their three types of preferences. Using the subjective expected utility model, only 20 of 64 subjects are not socially agnostic and, among them eleven are equity-preferring.²⁷

	EU: $\alpha = 1$		Subjective utility: $\alpha > 0$		
	γ value		α	γ value	
	VOI	Risk		VOI	Risk
Mean	1.61	1.16	0.53	3.81	2.77
Variance	4.46	3.34	0.25	40.25	8.69
Median	1.16	1.04	0.46	1.82	1.80
p25	0.72	0.73	0.09	1.15	1.09
p75	1.67	1.49	0.76	4.20	3.17

Estimates using nonlinear least squares (Heufer et al., 2018).

Table 5: Estimates of γ and α for both treatments (Table 5, Heufer et al. (2018))

While Schildberg-Hörisch (2010), Frignani and Ponti (2012) and Heufer et al. (2018) aim at eliciting distributive preferences of involved decision-makers under a veil of ignorance, Hong et al. (2015) use a similar methodology to identify distributive preferences of people put in a position of a social planner. Subjects had to choose an allocation between two other subjects given that $z = qy_1 + y_2$. They face twenty different budget sets varying in z and q . Their results might be

²⁵Note that subjects are fairly rational: there are only 3 violations of GARP in risk decision problems and 4 in decisions under the veil of ignorance. The mean value of the critical cost efficiency index (CCEI) is equal to 0.95.

²⁶This is the author's interpretation. However, it seems that in an allocation task under VOI, α can be interpreted as the relative weight one puts on the low payoff (the poor). If this is true, the results seem strange since the estimated mean value of α is about 0.5, suggesting that the rich gets a larger weight. Maybe α is not so easy to interpret, in that it captures both the subjective probabilities ratio and the relative weight of the poor.

²⁷For the other subjects, we can not reject that they are socially agnostic. This inconclusive result may be due to noise.

summarized in two points. First, as expected, the average share of the poorest decreases with q but the subjects were less sensitive to a change in larger prices. An increase in q from 1 to 5 decreases B 's share from 0.47 to 0.22, a further increase from 5 to 10 decreases the share to 0.17.²⁸ Second, they estimate a CES-utility function:²⁹

$$V(y_1, y_2) = [\alpha(y_1)^\rho + (1 - \alpha)(y_2)^\rho]^\frac{1}{\rho} \quad (7)$$

Figure 8 reports the distribution of the estimates of α and of the elasticity of substitution, i.e., $\sigma = \frac{1}{\rho-1}$. The median value of σ is found to be equal to -1.20 and σ is lower than -1 for a large part of subjects, that is, they value efficiency more than equity. Only about 20% of them have a preference for equity, i.e., σ higher than -1 .

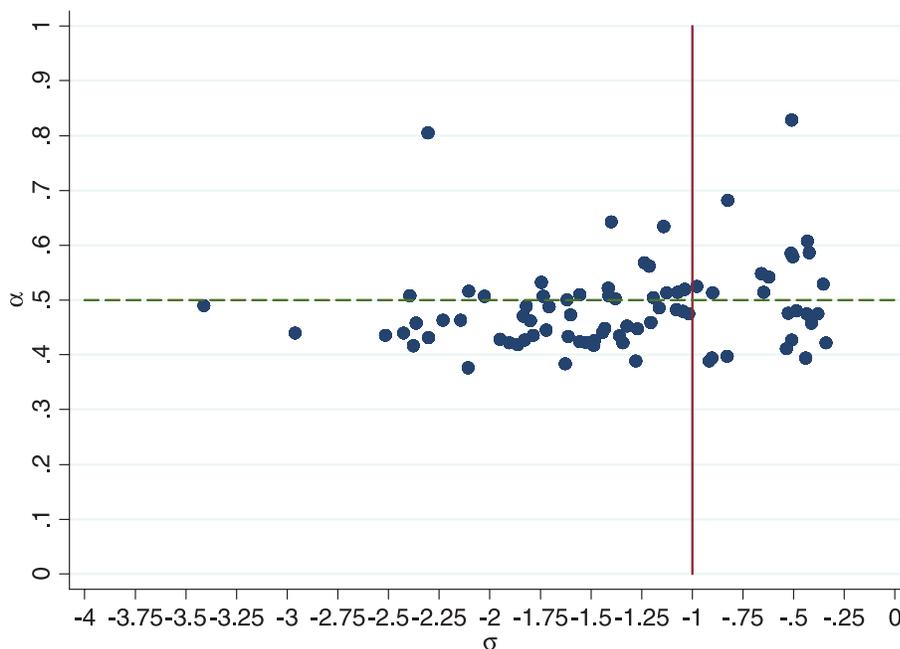


Figure 8: Estimates of σ and α (Figure 6, Hong et al., 2015)

This experimental research offers valuable insights, even though the case of two-person groups might be considered as restrictive. The extension of the framework to a setting with more than two stakeholders. raises a rather tricky theoretical question. Typically the utility functions that are assumed in this research are representations of a social ordering respecting some desirable properties such as the principle of progressive transfer. In a two-person environment, this may be acceptable, but when there are more than two stakeholders, previous research, including the experimental and

²⁸However, the use of elasticities may provide a slight different picture: One obtains arc elasticities of -0.54 and -0.38 and discrete elasticities of -0.13 and -0.23 .

²⁹The authors found that 56.9% of the subjects made choices that satisfied GARP but only 5% of them present a CCEI score lower than 0.97, that is almost all subjects are almost GARP-consistent.

questionnaire studies described in sections 4.2 and 3.2 respectively, has shown that this principle is far from being generally supported.

4.4 Ex ante versus ex post views

When we try to integrate risk and uncertainty in the normative evaluation of social situations, one of the most important questions is to determine whether we should adopt an ex ante or an ex post perspective. Within the prioritarian framework, the former approach asserts that \mathbf{v} is preferred to \mathbf{w} if the sum of individual's transformed expected well-being is larger in \mathbf{v} than in \mathbf{w} while, according to the latter, the comparison is based on the expected sum of individual's transformed well-being. In an Atkinson-setting with n individuals and S states, the two approaches can be written formally:

$$\begin{aligned}
 W^{EAP} &= \sum_{i=1}^n g \left(\sum_{s=1}^S \pi_s u(z_{is}) \right) \\
 W^{EPP} &= \sum_{s=1}^S \pi_s \sum_{i=1}^n g(u(z_{is})) \\
 &\text{with } f(u) = \frac{u^{1-\gamma}}{1-\gamma}
 \end{aligned}$$

where z_{is} is the payoff of i in state s , π_s the probability of state s , $u(\cdot)$ an individual utility function (with $u' > 0$ and $u'' < 0$ in case of risk aversion) and $g(\cdot)$ a transformation function (with $g' > 0$ and $g'' < 0$ in case of well-being inequality aversion). The ex ante perspective gives priority to the ex ante Pareto principle but violates a Dominance axiom, asserting that \mathbf{v} should be preferred to \mathbf{w} if \mathbf{v} is better than \mathbf{w} for *every* state of nature. On the contrary, the ex post approach respects Dominance but violates ex ante Pareto.

The papers that we report in this section aim at evaluating whether people adopt an ex ante or ex post view. Rohde and Rohde (2015) propose an experiment in which subjects in a position of a social planner have to rank different ten-stakeholders allocations of risks in a series of pairwise comparisons. The risky allocations to compare are the following:

- Independent lotteries (**I**): Each stakeholder independently faces the lottery $(p : \bar{y}, 1 - p : \underline{y})$, i.e., each stakeholder can be in one of the two states (she receives \bar{y} or \underline{y}).
- Common Lottery (**CL**): All the stakeholders face the common lottery $(p : \bar{y}, 1 - p : \underline{y})$, i.e., there are only two social states (all receive \bar{y} or all receive \underline{y}).
- Random Distribution (**R**): Five randomly drawn stakeholders receive \bar{y} and the others receive \underline{y} .
- IDbased Distribution (**Id**): Five stakeholders receive \bar{y} and the others receive \underline{y} based on the student ID-number.

- Common Outcome (**CO**): All receive $(\bar{y} + \underline{y})/2$.

In the experiment, three sets of parameters are tested: $(p, \bar{y}, \underline{y}) = (0.5, 10, 5)$, $(0.2, 11.5, 6.5)$ and $(0.7, 9, 3)$.

Let us focus on the two first risk allocations. The authors describe the choice between them as follows. If one seeks to minimize the *expected* standard deviation of incomes (‘pure’ ex post inequality minimizing choice), then CL should be preferred. If the objective is to minimize the standard deviation of expected incomes (ex ante inequality minimizing choice), then one should be indifferent between CL and I. However note that collective risk is smaller for the independent lotteries than for the common lottery. Thus if a social planner prefers CL to I, she reveals to have a strong preference for ex post equality (despite the increase in the collective risk) or to be collective-risk seeking. The overall ranking allows the authors to disentangle attitudes with respect to ex ante and ex post inequality and with respect to individual and collective risk.

The experiment was run at the Erasmus University of Rotterdam with 55 subjects. The aggregate ranking turns out to be: **CO** \sim **I** \succsim **CL** \succ **R** \succ **Id**, and a large majority of subjects are consistent with this overall ranking. To better understand this result, the authors estimate a logit econometric model in which the value of an allocation depends on ex ante and ex post inequality and individual and collective risk (both measured by the variance of payoffs). The subjects are found to be ex ante inequality averse, individual risk averse, ex post inequality seeking and collective risk seeking.

Another relevant experiment is conducted by Cettolin and Riedl (2016). Their subjects play the role of a social planner and are asked to allocate an amount of money, say z , between two *other* stakeholders, denoted C and U : $z = z_U + z_C$. The amount sent to C is certain, i.e., the amount of money allocated to C by the social planner is the one C receives. The amount received by U , denoted \tilde{z}_U , is not certain: U will receive $k^H z_U$ with probability p and $k^L z_U$ with probability $(1-p)$, where k^H and k^L are two multipliers such that $k^H > 1$ and $1 > k^L \geq 0$. The constraint imposed in the experiment is that the amount allocated to U is the expected payoff of U (then the expected total payoff remains constant). Formally:

$$z_U = E(\tilde{z}_U) = pk^H z_U + (1-p)k^L z_U.$$

The amount of z is determined in a preliminary production phase (in which the productivity levels are similar between the two stakeholders).³⁰ Five types of allocation problems (i.e., sets of parameters) are displayed in Table 6.

The authors test whether the choices made by their subjects can be rationalized by the following views of justice:

- Equality of expected payoffs (z_U should be always equal to 8);

³⁰The subjects performed a very simple ‘slider task’. Each group of stakeholders could earn a maximum of 16 euros and almost all groups earned 16 euros.

Allocation problem	Value of parameters			Amount given to U	
	p	k^H	k^L	Mean	Std dev
1-Certainty	1	1	.	8.39	1.59
2-Risk	0.5	1.5	0.5	8.50	1.43
3-Risk	0.8	1.25	0	8.13	1.69
4-Risk	0.5	2	0	7.33	2.10
5-Risk	0.2	5	0	5.10	2.84

Table 6: Amount given to U (Table 3, Cettolin and Riedl, 2016)

- Equality of expected utilities when utility is defined as $u(z) = z^\alpha$;
- Ex post equality when equality is measured by the absolute difference between payoffs;
- Utilitarianism, again with the utility function $u(z) = z^\alpha$.³¹

Table 7 reports the optimal choices according to the last three views (in the three last columns) given that $\alpha = 0.72$.³² The first columns of Table 7 report some predictions of choices of z_U according to ex ante and ex post prioritarianism (EAP and EPP) for different level of γ .³³ First, the EAP and EPP optimal choices are quite similar in a low risk situation but can be better distinguished when riskiness is larger. Second, due to the specific structure of the decision problem and the values of the parameters, a utilitarian planner and an ex post prioritarian one will choose the same level of z_U . Third, ex post equality and ex post prioritarianism also have similar predictions.

The authors run the experiment with a pool of 90 subjects of the Maastricht University School of Business and Economics. Their results are quite different from the ones obtained by Rohde and Rohde (2015). The two last columns of Table 6 indicate the mean value and standard deviation of the amount sent to U . First, the amount sent to U is decreasing with riskiness for rather high level

³¹The optimal choices are then:

$$\begin{aligned} \text{Equality of expected utilities: } z_U^{EEU} &= \frac{z}{K^{\frac{1}{\alpha}} + 1} \quad \text{with } K = p(k^H)^\alpha + (1-p)(k^L)^\alpha \\ \text{Ex post equality: } z_U^{EPE} &= \frac{1}{k^H + 1} z \quad \text{if } k^H > \frac{1-p}{p} \quad (\text{always true in the experiment}) \\ \text{Utilitarianism: } z_U^U &= \frac{z}{K^{\frac{1}{\alpha-1}} + 1} \end{aligned}$$

³²The authors elicited in a pre-experimental task subject's belief about other's risk preference. The subject's mean level of α is 0.72.

³³The optimal choices of ex ante and ex post prioritarianism are:

$$\begin{aligned} z_U^{EAP} &= \frac{z}{K^{\frac{1-\gamma}{\alpha(1-\gamma)-1}} + 1} \quad \text{with } K = p(k^H)^\alpha + (1-p)(k^L)^\alpha \\ z_U^{EPP} &= \frac{z}{M^{\frac{1}{\alpha(1-\gamma)-1}} + 1} \quad \text{with } M = p(k^H)^{\alpha(1-\gamma)} + (1-p)(k^L)^{\alpha(1-\gamma)} \end{aligned}$$

	EAP	EPP	EAP	EPP	EAP	EPP	Utilitarian	EU equality	Ex Post equality
	$\gamma = 0.1$		$\gamma = 0.5$		$\gamma = 0.9$		$\gamma = 0$		
2-Risk	7.72	7.64	7.91	7.80	7.99	7.96	7.61	8.15	6.4
3-Risk	7.36	7.11	7.80	7.11	7.97	7.11	7.11	8.35	7.11
4-Risk	6.05	5.33	7.39	5.33	7.92	5.33	5.33	9.07	5.33
5-Risk	3.84	2.67	6.61	2.67	7.81	2.67	2.67	10.42	2.67

Table 7: Some predictions - choice of z_U

of riskiness. There is no clear difference in the mean value of z_U between certainty and 2-Risk/3-Risk problems while it seems significantly lower in 4-Risk/5-Risk problems. Second, the dispersion becomes larger with riskiness suggesting that there exists a strong heterogeneity across subjects. The authors then estimate for each social planner the view of justice that fits best her choices.³⁴ They find that 43% of the subjects equalize expected payoffs, 36% act as if they are utilitarian, 14% seek ex post equality, and 7% seek equality of expected utilities.

Finally, Andreoni et al. (2020) study the problem of allocating lottery tickets. Twenty lottery tickets have to be allocated between two poor (and equally-deserving) Kenyan households, denoted A and B . Ten of these tickets are randomly allocated by a computer and the problem of a social planner is to allocate the 10 remaining ones. Next one ticket is randomly drawn and the household to which the ticket is assigned receives \$10. Three classes of decision problems are studied. In a first class of problems (called *ex ante decision*), the decision of the social planner comes after learning the allocation of the computer. In a second class (called *ex post decision*), the decision is made *after* learning whether the winning ticket is one of the computer's tickets or one of hers. A third one (called *ex ante decision with revision*) is similar to the *ex ante decision* frame except that, after her decision, the social planner is given the information whether the winning ticket is one of the computer's tickets or one of hers and she can change her initial decision.

Suppose that the computer assigned 6 tickets to A . Let us first focus on an ex ante prioritarian social planner. The pattern of choices should be as follows: (i) she allocates 4 tickets to A in an *ex ante decision* frame; (ii) 5 tickets to A in an *ex post decision* frame; and (iii) revises her allocation in *ex ante decision with revision*. Andreoni et al. (2020) consider such a pattern of choices as time-inconsistent since the social planner fails to stick to his initial plan.³⁵ Then a fair ex ante consequentialist should commit herself against any change in choices. On the contrary, if one seeks to minimize ex post inequality of *earnings*, all the allocations of tickets are equivalent in terms of desirability since the same allocation, i.e., $(0, 10)$, is reached.³⁶

³⁴As for the level of α , the authors use the social planner's belief about risk aversion of others, elicited in a pre-experimental task.

³⁵However one could also think that it is rational to change one's mind after receiving new (and relevant) information.

³⁶Note that Andreoni et al. (2020) classifies as ex post equalizing choice a choice that gives five tickets to both potential recipients.

Based on a pool of subjects of the University of California, San Diego, the authors report that most subjects equalize the chances of having the winning tickets in *ex ante decision* (more than 60% of choices) and *ex post decision* (about 45% of choices) settings. Moreover, when subjects have the opportunity to change their allocation after knowing the computer's allocations of its tickets (*ex ante decision with revision*), they tend to switch to the ex post equalizing choice, i.e., five tickets for both households.³⁷ Moreover, in an additional treatment in which subjects are free to make a contingent plan, they tend to plan to switch an ex ante equalizing choice to an ex post equalizing choice.

The authors also present a variant of this experiment in which subjects are invited to allocate \$10 between the two households in the same frames, given that there is a 50% probability that her choice is selected for payment and a 50% probability that the computer's choice is selected. Note that, according to an ex post inequality-averse view, one should always choose the allocation (5, 5). As for a social planner with a sufficiently large ex ante inequality aversion, she will choose an allocation that permits to equalize the expected utilities and may therefore change her choice once it is revealed whether it is the computer or her allocation that is selected for payment. The subjects are found to have a preference for equalizing expected utilities in *ex ante decision* (about 50%) even though the choice of equalizing the ex post distribution of payoffs represents a substantial part of choices (about 20%). When it is revealed that the subject's allocation is used for payment, subjects revise their choice and opt for an egalitarian allocation, i.e., \$5 for each household (about 50% of choices).

All in all, according to the author's interpretation, this result contradicts fair (ex ante or ex post) consequentialism and indicates that subjects exhibit a time-consistent preference for applying a naive deontological principle (here equality).

These three papers give very different pictures of what could be the people's ethical preferences in a risky or uncertain world. It is difficult to reconcile all these findings without any further research. Andreoni et al. (2020) stress an important (methodological) issue. Subjects probably have some difficulties to grasp the decision of allocating uncertain resources (or chances) among people. The decision framing might thus play an important role. For instance, in the Rohde and Rohde (2015) experiment, allocations are displayed in such a way that subjects consider the problem in an ex ante perspective and do not pay attention to the possible ex post inequalities. As Andreoni et al. (2020) emphasize, lowering the cognitive overload associated with such a task may be an important challenge if we want to understand people's preferences. This brings us back to one of the basic insights of this chapter. It is not easy to "test" sophisticated ethical theories, such as prioritarianism, on the basis of the answers and/or the behavior of unsophisticated agents. On the other hand, in the real world people also take their decisions in framed situations. Understanding how the frame influences their distributive preferences may therefore be an interesting research

³⁷The findings of additional treatments indicate that such a preference is strict and can not be explained by the resolution of indifference. That is, their choices are not consistent with ex post consequentialism.

question. Finally, it can be pointed out that these papers aim at studying what choices made by subjects in an allocation problem reveal about their underlying subjective social welfare function. Instead, it would be interesting to test directly the empirical support for basic principles such as the dominance axiom and the ex ante Pareto principle.

4.5 A broader set of questions: welfarism, desert and taxation

We have described in section 3.5 how questionnaire studies have found over and over again that the source of income differences is a crucial factor in distributive judgments. This finding is corroborated by the growing literature in experimental economics studying the role of effort and choices influencing the initial inequalities or the total income (Frohlich and Oppenheimer, 1990, Konow, 2000, Dickinson and Tiefenthaler, 2002, Cappelen et al., 2013, Durante et al., 2014, Almas et al., 2010).³⁸ The method used in these papers consists in adding a *production phase* to the game: that might be an effort task like a quiz (Dickinson and Tiefenthaler, 2002, Durante et al., 2014), a computer game (Tetris in Durante et al., 2014), preparing letters for mailing (Konow, 2000) or risk-taking choice (Cappelen et al., 2013, Akbas et al., 2019, Mollerstrom et al., 2015). In all these studies, people are found to be less willing to redistribute when income differences are due to effort.

In a recent paper, Almas et al. (2020) report an experiment in which a social planner has the opportunity to redistribute income between two workers, that are recruited from the platform Amazon Mechanical Turk to complete small tasks. They are paired by groups. One worker is initially assigned 6 USD and the other 0 USD for completing the task. In the *luck treatment*, the assignment is randomly determined; in the *merit treatment*, the assignment is made according to the productivity level. The social planner has then the choice between seven allocations, from (6,0) to (0,6), and there is no cost of redistribution.³⁹ The experiment was run online with Norwegian and US subjects. When income differences reflect luck, about 80% of Norwegian social planners and 55% of US ones choose the equal distribution (3,3). In the merit treatment, social planners accept a higher level of inequality. Allocation (4,2) is the preferred option in Norway (more than 40% of Norwegian social planners) and (3,3) is chosen by 35% of them. Among the US social planners, allocations (4,2) and (6,0) are the most attractive (with about 40% of US social planners). This echoes the findings of intercultural differences that were reported in section 2.

Some papers tackle the role of brute and option luck. In an experiment conducted by Cappelen et al. (2013), social planners can distribute a total amount of money between two subjects given that the total amount is determined by a choice between lotteries (a risky lottery versus a safe lottery). The authors found that social planners are willing to compensate inequality due to luck between two risk takers, but do not make transfers from a safer to an unlucky risk-taker (see also Akbas et

³⁸We report here the experiments that aim to identify the ethical preferences of social planner. There is also a growing literature studying how the sources of inequality affect ethical preferences or preferences for redistribution. See Cappelen et al. (2007), Cappelen et al. (2010), Almas et al. (2010), Balafoutas et al. (2013), Durante et al. (2014).

³⁹They also ran an additional treatment where transferring earnings to the poorest worker is costly.

al., 2019). Further, responsibility seems to be considered by social planners in a complex way. In an experiment where subjects face risk but have the opportunity to insure against it, Mollerstrom et al. (2015) report that social planners are less willing to compensate inequality due to bad brute luck when stakeholders do not make effort to control randomness. As mentioned already in section 3.5, this almost universal finding that the source of income differences matters, raises an important challenge for prioritarian welfarism.

Few experimental papers study ethical preferences over taxation. In Durante et al. (2014), subjects are invited to choose a tax and redistribution scheme that consists in a proportional tax rate with a demogrant. The authors study preferences for redistribution in different contexts: position (involved behind a veil of ignorance or not), efficiency cost, sources of income. However, they use taxation as a frame to elicit ethical and social preferences but do not study whether people's ethical preferences are different in the specific context of taxation. In a recent paper, Charité et al. (2015) show that subjects in a position of redistributing unequal and unearned initial endowments between two recipients do not make the same choice when the recipients know their initial endowment or not. It is found that allocators tend to be less willing to redistribute when initial endowments are known. The authors explain their result by the fact that allocators anticipate loss aversion, initial endowment being the reference point. Again, this result seems to contradict prioritarian consequentialism according to which pre-tax incomes should not matter for the evaluation of the taxation scheme.

5 “Testing” prioritarianism?

Let us now take stock and ask how relevant this empirical work is to better understand the popular acceptance of prioritarianism, or to get an estimate of the value of the crucial parameters, as perceived by the population. Let us start with the questionnaire studies. Traditionally, economists have been rather reluctant to take unincentivized answers by respondents seriously, but this attitude is recently changing. And rightly so. Answers on questionnaires should indeed not be trusted if we want to predict behavior, but they are very useful to get a better insight in ethical preferences. Ethical preferences are only one of the determinants of behavior, but they are worth investigating for their own sake, e.g. to check whether a given normative framework, such as prioritarianism, is popular among the population.

We have mentioned already the main limitation of the questionnaire studies on prioritarianism: the respondents are asked to evaluate income (or health) distributions, not distributions of well-being. This makes it impossible to interpret the results directly within a prioritarian framework. More specifically, it is not possible to rigorously distinguish utilitarianism and prioritarianism, as a utilitarian would also advocate income redistribution under the assumption of concave individual well-being functions. Analogously, the empirical estimates of the degree of inequality aversion relate to income or health levels, or to income-related health inequalities, not to individual well-being

levels, and therefore do not yield an estimate of the inequality aversion η in the general prioritarian social welfare function analyzed in this book.

This does not mean, of course, that the results are not informative. The opposition to the basic axioms underlying prioritarianism, such as the Pigou-Dalton transfer principle or the axiom of monotonicity, remains relevant even when we take into account that they are formulated in the space of incomes. The rejection of monotonicity suggests that income externalities play some role in the ethical preferences of the respondents. Rejection of the transfer principle is even more fundamental, as one would certainly expect this principle to be satisfied within a prioritarian framework. Here also externalities may play a role. As we have seen, the transfer principle is most often rejected when the income transfer tends to “isolate” the poor, which suggests that respondents imagine that the growing income gap between the poorest and the second poorest decreases the well-being of the poorest even if their income does not change. More research is needed to better understand the rationale for the rejection of the transfer principle in the space of incomes.

The partial results on inequality aversion may also be useful inputs if the aim is to get a better insight into the structure of the prioritarian social welfare function. Sticking as before to the Atkinson framework, we can write this function as

$$S = \frac{1}{1-\gamma} \sum_i w(y_i, H_i)^{1-\gamma}, \quad (8)$$

which can be directly compared to the simpler forms (1)-(3). Studies such as those described in section 3.3 (yielding estimates for e and ε) should be helpful to think about the specification of $w(\cdot)$ and to estimate the parameter γ . In fact, in a von Neumann-Morgenstern setting this work could be combined with other empirical work aiming at estimating the degree of risk aversion in the function $w(\cdot)$. Interesting examples of this approach are Johansson-Stenman et al. (2002) and Carlsson et al. (2005), but they focus on relative incomes and interpret inequality as affecting individual well-being rather than as a normative parameter which does not necessarily depend on individual self-interested preferences.

The pros and cons of laboratory experiments to some extent mirror the pros and cons of questionnaire studies. Basically, they investigate how individuals in a lab share wealth within a small group. To be relevant for a discussion of prioritarianism, we need to make the crucial assumption that the choices made by subjects reveal their true ethical preferences. In a significant part of these experiments, subjects in a position of a social planner are asked to share an amount of money between other subjects. Some could argue that such experiments do not provide any informative insights because of the lack of incentives. This criticism is similar to the traditional criticism on questionnaire studies. In another set of experiments, participants are put behind a veil of ignorance. Although this procedure is suggested by (a part of) the philosophical literature, one can argue that participants behind a VOI are likely to reveal their risk preferences rather than their ethical ones. Second, even if the lab experiments correctly reveal the participant’s ethical preferences, these can

only be the preferences on (re)distribution within a small group. Does this say anything about their ethical preferences in a large society? One might think that these two contexts (a small group of students in a lab and a large and diverse society) imply different types of ethical reasoning. First, it is difficult to capture in a lab all relevant dimensions that may play a role in larger groups. For instance, transferring wealth between two persons might be differently evaluated depending on the presence of others. This echoes the findings about Pigou-Dalton in the questionnaires. Second, a laboratory is a small artificial society in which the norm is generally to earn as much money as possible.⁴⁰ A society is a set of people sharing values, norms, friendship, solidarity and identity that we fail to capture in the lab. For instance, persons may attach more importance to efficiency than to equity when sharing a resource in the lab than they would outside – even though they are not involved.

Yet, despite these problems, lab experiments remain very useful. First, we believe that one can be confident that what is revealed in a lab are the participant’s ethical preferences. Assume for instance that participant’s preferences are represented by a utility function such as $v(x) = u(x_i) + \alpha w(x)$ where $x_i = (x_1, \dots, x_i, \dots)$ is the social state, x_j the payment of individual j , u is the self-centered utility function, w is the social welfare function and α is the weight put on the social welfare relative to her own self-centered utility. Then, in a position of a non-involved spectator, x_i is constant and her choice of a distribution of payoffs should reveal her ethical ranking. Moreover, the empirical work tend to show that, at least in the experiments using the spectator framework, participants reveal their ethical preferences rather than making only random choices (e.g., Cettolin and Riedl, 2016).

Second, eliciting ethical preferences in the field is almost impossible. It is very rare that people have to allocate resources within a group and, if this is the case (e.g., a manager who sets wages of her employees), decisions are not determined only by ethical preferences. Lab-experimental work may then helpfully complement questionnaire-experimental studies. In particular, it might inform whether stated ethical preferences are consistent with allocative choices when there is a *real* stake. Moreover, placing individuals in the position of allocating a resource within a small group can be a good way to remove any external disturbance, such as those described in section 2. A prime example is the biased perception of the actual situation. As we argued earlier, in a reflective equilibrium setting, we confront general principles with “relevant normative facts”, and behavior in the lab can certainly be seen as normatively relevant information.

Finally, a significant part of the empirical studies we reported put the respondent in a position of choosing a wealth/health distribution behind the Harsanyi or Rawlsian veil of ignorance. There may indeed be doubts whether the veil of ignorance is the appropriate standpoint for determining what is fair or unfair. Harsanyi has argued that an agent behind the veil of ignorance should behave as an utilitarian social planner, reinforcing the view that it is not *rational* to be prioritarian. However, some authors argue that such a result is questionable on the ground that it is not obvious

⁴⁰The relative acceptance of the Pareto axiom in lab experiments might be explained by such a norm.

that his representation theorem justifies utilitarianism (see Roemer (1996) for a summary of the debate).⁴¹ Moreover, it is not clear that society's preferences over social states should be derived from preferences behind a veil of ignorance (see Sen, 2009, for a deeper discussion). The finding, both in the questionnaire studies and in the lab experiment, that subjects in a role of external observer (or social planner) take different decisions than behind a veil of ignorance, is certainly relevant in this whole discussion. Moreover, even if we do believe that the uninvolved spectator is the *most* appropriate view for judging the fairness of situation, decisions made by individuals facing lotteries of social states might still contribute to our understanding of people's ethical preferences.

6 Conclusion

The results summarized in this chapter are useful to assess the political feasibility of prioritarian policy evaluation in a democratic society. Another question is whether they can be used directly to guide normative tax analysis (as is suggested, e.g., by Weinzierl, 2014). This seems premature, as most of the questions raised in the representative questionnaire studies are rather simplistic, and the setting of the behavioral experiments is highly stylized. We first have to understand better the underlying structure of the ethical attitudes and their robustness. As we have seen, individuals change their views when they receive more information. Moreover, while we sometimes have seen strong majorities on some issues, there is almost never a full social consensus. A large degree of preference heterogeneity is found both in the survey studies and in the lab. Society therefore still has to devise an acceptable procedure to aggregate conflicting attitudes. It is tempting for social researchers to just take the average (or the median, or the majority) answers to specific questions, but it remains to be investigated if and under which circumstances these simple procedures yield socially desirable outcomes.

The relevance of these results from a normative point of view raises similar questions. The fact that welfarism in general, and more specifically prioritarianism, do not seem to be the dominant ethical views among the population of our societies is not an argument against them as a chosen ethical framework. It still raises two challenges, however. From a sociopolitical point of view, it gives us a good idea of the difficult task that lies ahead of us if we want to foster sufficient democratic support for the prioritarian policies that are described in the other chapters of this volume. From an academic point of view, it invites us to confront explicitly the deviating opinions. Is the Pigou-Dalton transfer principle indeed a crucial building block of our social welfare function - or is there some truth in the widely spread view that it becomes problematic in a society with more than two individuals? Does the fact that preferences of respondents behind a veil of ignorance are different from their preferences as a social planner indicate that the veil of ignorance is a misleading

⁴¹Indeed, Harsanyi (1955) shows that the sum of individual vNM utilities is the only function that represents a social ordering over a lottery of social states when the social preference and individual's preference over social states respect the vNM axioms. It is not clear that the summation can be interpreted as the sum of individual 'utilities'.

device to recover social welfare judgments? Are ideas like “equal sacrifice” or “progressive taxation as a value in itself” ethically attractive? If they are, should they be integrated in prioritarian thinking? Or should they be seen as independent principles? But then, how to think about the trade-off between respecting these principles and optimizing a social welfare function? How to handle the dominant view that individuals should be held responsible for part of their outcomes? Is the normative definition of well-being that is explored in theories of fair allocation an attractive way out, or does this approach remove us too far from the basic intuition that individual well-being is not only normative, but has necessarily some underlying psychological basis? The empirical research on distributive preferences cannot answer these questions, but it has the merit to put them on the table.

References

- Abasolo, I. and Tsuchiya, A. (2004). “Exploring social welfare functions and violations of monotonicity: an example from inequalities in health.” *Journal of Health Economics*, 23: 313-329.
- Aguiar, F., Becker, A. and Miller, L. (2013). “Whose impartiality? An experimental study of veiled stakeholders, involved spectators and detached observers.” *Economics and Philosophy*, 29: 155-174.
- Akbas, M., Ariely, D. and Yuksel, S. (2019). “When is inequality fair? An experiment on the effect of procedural justice and agency.” *Journal of Economic Behavior and Organization*, 161: 114-127.
- Alesina, A. and Angeletos, G.-M. (2005). “Fairness and redistribution.” *American Economic Review*, 95: 960-980.
- Alesina, A. and Giuliano, P. (2011). “Preferences for redistribution.” In Benhabib, J., Bisin, A. and Jackson, M., eds., *Handbook of Social Economics, Vol. 1A*. New York: Elsevier, pp. 93-131.
- Alesina, A., Stantcheva, S. and Teso, E. (2018). “Intergenerational mobility and support for redistribution.” *American Economic Review*, 108: 521-554.
- Almas, I., Cappelen, A. W., Sorensen, E., and Tungodden, B. (2010). “Fairness and the development of inequality acceptance.” *Science*, 328: 1176-1178.
- Almas, I., Cappelen, A. W. and Tungodden, B. (2020). “Cutthroat capitalism versus cuddly socialism: Are Americans more meritocratic and efficiency-seeking than Scandinavians?” *Journal of Political Economy*, forthcoming.
- Amiel, Y. and Cowell, F. (1994). “Monotonicity, dominance and the Pareto principle.” *Economics Letters*, 45: 447-450.
- Amiel, Y. and Cowell, F. (1999). *Thinking about inequality*. Cambridge: Cambridge University Press.
- Amiel, Y., Cowell, F. and Gaertner, W. (2009). “To be or not to be involved: a questionnaire-experimental view on Harsanyi’s utilitarian ethics.” *Social Choice and Welfare*, 32: 253-274.
- Amiel, Y., Cowell, F. and Gaertner, W. (2012). “Distributional orderings: an approach with seven flavours.” *Theory and Decision*, 73: 381-399.
- Amiel, Y., Creedy, J. and Hurn, D. (1999). “Attitudes towards inequality.” *Scandinavian Journal of Economics*, 101: 83-96.
- Andreoni, J., Aydin, D., Barton, B., Bernheim, B. D., and Naecker, J. (2020). “When fair isn’t fair: Understanding choice reversals involving social preferences.” *Journal of Political Economy*, 128: 1673-1711.
- Andreoni, J. and Miller, J. (2002). “Giving according to garp: An experimental test of the consistency of preferences for altruism.” *Econometrica*, 70: 737-753.
- Balafoutas, L., Kocher, M. G., Putterman, L. and Sutter, M. (2013). “Equality, equity and incentives: an experiment.” *European Economic Review*, 60: 32-51.
- Ballano, C. and Ruiz-Castillo, J. (1993). “Searching by questionnaire for the meaning of income inequality.” *Revista Espanola de Economia*, 10: 233-259.
- Becker, A. and Miller, L. M. (2009). “Promoting justice by treating people unequally: an experimental study.” *Experimental Economics*, 12: 437-449.
- Beckman, S. R., Formby, J. P., Smith, W. J. and Zheng, B. (2002). “Envy, malice and pareto efficiency: an experimental examination.” *Social Choice and Welfare*, 19: 349-367.
- Benabou, R. and Ok, E. (2001). “Social mobility and the demand for redistribution: the POUM hypothesis.” *Quarterly Journal of Economics*, 116: 447-487.
- Benabou, R. and Tirole, J. (2006). “Belief in a just world and redistributive politics.” *Quarterly Journal of Economics*, 121: 699-746.

- Bernasconi, M. (2002). "How should income be divided? Questionnaire evidence from the theory of "impartial preferences". *Journal of Economics* Supplement, 9: 163-195.
- Bleichrodt, H., Rohde, K. and Van Ourti, T. (2012). "An experimental test of the concentration index." *Journal of Health Economics*, 31: 86-98.
- Bosmans, K. and Schokkaert, E. (2004). "Social welfare, the veil of ignorance and purely individual risk: an empirical examination." *Research on Economic Inequality*, 11: 85-114.
- Cappelen, A. W., Hole, A. D., Sorensen, E. and Tungodden, B. (2007). "The pluralism of fairness ideals: an experimental approach." *American Economic Review*, 97: 818-827.
- Cappelen, A. W., Konow, J., Sorensen, E. and Tungodden, B. (2013). "Just luck: An experimental study of risk-taking and fairness." *American Economic Review*, 103: 1398-1413.
- Cappelen, A. W., Sorensen, E. and Tungodden, B. (2010). "Responsibility for what? Fairness and individual responsibility." *European Economic Review*, 54: 429-441.
- Carlsson, F., Daruvala, D. and Johansson-Stenman, O. (2005). "Are people inequality averse or just risk averse?" *Economica*, 72: 375-396.
- Cettolin, E. and Riedl, A. (2016). "Justice under uncertainty." *Management Science*, 63: 3739-3759.
- Charité, J., Fisman, R., and Kuziemko, I. (2015). "Reference points and redistributive preferences: Experimental evidence." *NBER Working Paper* No. 21009.
- Charness, G. and Rabin, M. (2002). "Understanding social preferences with simple tests." *Quarterly Journal of Economics*, 117: 817-869.
- Clark, A. and D'Ambrosio, C. (2015). "Attitudes to income inequality: experimental and survey evidence." In Atkinson, A. and Bourguignon, F., eds., *Handbook of Income Distribution, Vol. 2A*. New York: Elsevier, pp. 1147-1208.
- Cookson, R., Ali, S., Tsuchiya, A. and Asaria, M. (2018). "E-learning and health inequality aversion: a questionnaire experiment." *Health Economics*, 27: 754-1771.
- Cooper, D. and Kagel, J. H. (2015). "Other-regarding preferences: A selective survey of experimental results." In Kagel, J. H. and Roth, A. E., eds., *Handbook of experimental economics, Vol. 2*. Princeton: Princeton University Press, pp. 217-89.
- Costa-Font, J. and Cowell, F. (2019). "Incorporating inequality aversion in health-care priority setting." *Social Justice Research*, 32: 172-185.
- Cruces, G., Perez-Truglia, R. and Tetaz, M. (2013). "Biased perceptions of income distribution and preferences for redistribution: evidence from a survey experiment." *Journal of Public Economics*, 98: 100-112.
- Dickinson, D. L. and Tiefenthaler, J. (2002). "What Is Fair? Experimental Evidence." *Southern Economic Journal*, 69: 414-428.
- Dolan, P. and Tsuchiya, A. (2011). "Determining the parameters in a social welfare function using stated preference data: an application to health." *Applied Economics*, 43: 2241-2250.
- Durante, R., Putterman, L., and van der Weele, J. (2014). "Preferences for redistribution and perception of fairness: an experimental study." *Journal of the European Economic Association*, 12: 1059-1086.
- Edlin, R., Tsuchiya, A. and Dolan, P. (2012). "Public preferences for responsibility versus public preferences for reducing inequalities." *Health Economics*, 21: 1416-1426.
- Engelmann, D. and Strobel, M. (2004). "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments." *American Economic Review*, 94: 857-869.

- Fehr, E., Naef, M., and Schmidt, K. M. (2006). "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: comment." *American Economic Review*, 96: 1912-1917.
- Fisman, R., Kariv, S., and Markovits, D. (2007). "Individual preferences for giving." *American Economic Review*, 97: 1858-1876.
- Fleurbaey, M. and Maniquet, F. (2018). "Optimal income taxation theory and principles of fairness." *Journal of Economic Literature*, 56: 1029-1079.
- Fleurbaey, M. and Maniquet, F. (2011). *A theory of fairness and social welfare*. Cambridge: Cambridge University Press.
- Fong, C. (2001). "Social preferences, self-interest, and the demand for redistribution." *Journal of Public Economics*, 82: 225-246.
- Frignani, N. and Ponti, G. (2012). "Risk versus social preferences under the veil of ignorance." *Economics Letters*, 116: 143-146.
- Frohlich, N. and Oppenheimer, J. A. (1990). "Choosing justice in experimental democracies with production." *American Political Science Review*, 84: 461-477.
- Frohlich, N., Oppenheimer, J. A., and Eavey, C. L. (1987a). "Choices of principles of distributive justice in experimental groups." *American Journal of Political Science*, 31: 606-636.
- Frohlich, N., Oppenheimer, J. A., and Eavey, C. L. (1987b). "Laboratory results on Rawls's distributive justice." *British Journal of Political Science*, 17: 1-21.
- Gaertner, W. and Schokkaert, E. (2012). *Empirical social choice*. Cambridge: Cambridge University Press.
- Giuliano, P. and Spilimbergo, A. (2014). "Growing up in a recession." *Review of Economic Studies*, 81: 787-817.
- Glejser, H., Gevers, L., Lambot, P. and Morales, J. (1977). "An econometric study of the variables determining inequality aversion among students." *European Economic Review*, 10: 173-188.
- Güth, W. and Kliemt, H. (2010). "What ethics can learn from experimental economics - if anything." *European Journal of Political Economy*, 26: 302-310.
- Harrison, E. and Seidl, C. (1994). "Acceptance of distributional axioms: experimental findings." In Eichhorn, W., ed., *Models and measurement of welfare and inequality*. Berlin: Springer-Verlag, pp. 67-99.
- Harsanyi, J. (1953). "Cardinal utility in welfare economics and the theory of risk-taking." *Journal of Political Economy*, 61: 434-435.
- Harsanyi, J. (1955). "Cardinal welfare, individualistic ethics and interpersonal comparisons of utility." *Journal of Political Economy*, 63: 309-321.
- Hausman, D. (2015). *Valuing health: well-being, freedom, and suffering*. Oxford: Oxford University Press.
- Herne, K. and Suojanen, M. (2004). "The role of information in choices over income distributions." *Journal of Conflict Resolution*, 48: 173-193.
- Heufer, J., Shachat, J., and Xu, Y. (2018). "Measuring tastes for equity and aggregate wealth behind the veil of ignorance." *Tinbergen Institute Discussion Paper*, TI 2018-087/I.
- Hong, H., Ding, J., and Yao, Y. (2015). "Individual social welfare preferences: an experimental study." *Journal of Behavioral and Experimental Economics*, 57: 89-97.
- Isaksson, A.-S. and Lindskog, A. (2009). "Preferences for redistribution - a country comparison of fairness judgements." *Journal of Economic Behavior and Organization*, 72: 884-902.

- Jancewicz, B. (2016). "Income inequalities: axioms of income inequality measures and people's perceptions." *Decyzje*, 5: 21-42.
- Johansson-Stenman, O., Carlsson, F. and Daruvala, D. (2002). "Measuring future grandparents' preferences for equality and relative standing." *Economic Journal*, 112: 363-383.
- Konow, J. (2000). "Fair shares: Accountability and cognitive dissonance in allocation decisions." *American Economic Review*, 90: 1072-1091.
- Konow, J. (2001). "Fair and square: the four sides of distributive justice." *Journal of Economic Behavior and Organization* 46: 137-164.
- Konow, J. (2003). "Which is the fairest one of all? A positive analysis of justice theories." *Journal of Economic Literature*, 41: 1188-1239.
- Kuhn, A. (2013). "Inequality Perceptions, Distributional Norms, and Redistributive Preferences in East and West Germany." *German Economic Review*, 14: 483-499.
- Kuziemko, I., Norton, M., Saez, E. and Stantcheva, S. (2015). "How elastic are preferences for redistribution? Evidence from randomized survey experiments." *American Economic Review*, 105: 1478-1508.
- Lal, A., Mohebi, M., Sweeney, R., Moodie, M., Peeters, A. and Carter, R. (2019). "Equity weights for socioeconomic position: two methods - survey of stated preferences and epidemiological data." *Value in Health*, 22: 247-253.
- Leibler, J., Zwack, L. and Levy, J. (2009). "Agreement with inequality axioms and perceptions of inequality among environmental justice and risk assessment professionals." *Health, Risk & Society*, 11: 55-69.
- Luce, R. and Raiffa, H. (1957). *Games and decisions*. New York: Wiley.
- Luttmer, E. and Singhal, M. (2011). "Culture, context, and the taste for redistribution." *American Economic Journal: Economic Policy*, 3: 157-179.
- McClelland, G. and Rohrbaugh, J. (1978). "Who accepts the Pareto axiom? The role of utility and equity in arbitration decisions." *Behavioral Science*, 23: 446-456.
- Miller, D. (1992). "Distributive justice: what the people think." *Ethics*, 102: 555-593.
- Mollerstrom, J., Reme, B.-A., and Sorensen, E. (2015). "Luck, choice and responsibility. An experimental study of fairness views." *Journal of Public Economics*, 131: 33-40.
- Muller, D. and Renes, S. (2020). "Fairness views and political preferences: evidence from a large and heterogeneous sample." *Social Choice and Welfare*, <https://doi.org/10.1007/s00355-020-01289-5>.
- Nord, E. (1995). "The person-trade-off approach to valuing health care programs." *Medical Decision Making* 15: 201-208.
- Ottersen, T., Maestad, O. and Norheim, O. (2014). "Lifetime QALY prioritarianism in priority setting: quantification of the inherent trade-off." *Cost Effectiveness and Resource Allocation* 12.
- Pinto-Prades, J.-L. and Abellán-Perpiñán (2005). "Measuring the health of populations: the veil of ignorance approach." *Health Economics*, 14: 69-82.
- Pirttilä, J. and Uusitalo, R. (2010). "A 'leaky bucket' in the real world: estimating inequality aversion using survey data." *Economica*, 77: 60-76.
- Rawls, J. (1958). "Justice as fairness." *Philosophical Review*, 67: 164-194.
- Rawls, J. (1971). *A theory of justice*. Cambridge, Mass.: Harvard University Press.
- Robson, M., Asaria, M., Cookson, R., Tsuchiya, A. and Ali, S. (2017). "Eliciting the level of health inequality aversion in England." *Health Economics*, 26: 1328-1334.
- Roemer, J. (1996). *Theories of distributive justice*. Cambridge, Mass.: Harvard University Press.

- Rohde, I. M. and Rohde, K. I. (2015). "Managing social risks - tradeoffs between risks and inequalities." *Journal of Risk and Uncertainty*, 51: 103-124.
- Saez, E. and Stantcheva, S. (2013). "Generalized social marginal welfare weights for optimal tax theory." *NBER: Working Paper* 18835.
- Schildberg-Hörisch, H. (2010). "Is the veil of ignorance only a concept about risk? An experiment." *Journal of Public Economics*, 94: 1062-1066.
- Schokkaert, E. and Truyts, T. (2017). "Preferences for redistribution and social structure." *Social Choice and Welfare*, 49: 545-576.
- Sen, A. (2009). *The idea of justice*. London: Allen Lane.
- Tarroux, B. (2015). "Comparing two-dimensional distributions: a questionnaire-experimental approach." *Social Choice and Welfare*, 44: 87-108.
- Tarroux, B. (2019). "The value of tax progressivity: evidence from survey experiments." *Journal of Public Economics* 179: Article 104068.
- Traub, S., Seidl, C., and Schmidt, U. (2009). "An experimental study on individual choice, social welfare, and social preferences." *European Economic Review*, 53: 385-400.
- Traub, S., Seidl, C., Schmidt, U., and Levati, M. V. (2005). "Friedman, Harsanyi, Rawls, Boulding - or somebody else? An experimental investigation of distributive justice." *Social Choice and Welfare*, 24: 283-309.
- Weinzierl, M. (2014). "The promise of positive optimal taxation: normative diversity and a role of equal sacrifice." *Journal of Public Economics*, 118: 128-142.
- Weinzierl, M. (2017). "Popular acceptance of inequality due to innate brute luck and support for classical benefit-based taxation." *Journal of Public Economics*, 155: 54-63.
- Weinzierl, M. (2017). "A welfarist role for nonwelfarist rules: an example with envy." *Harvard Business School: Working Paper* 17-021.
- Yaari, M. and Bar-Hillel, M. (1984). "On dividing justly." *Social Choice and Welfare*, 1: 1-24.
- Yamamura, E. (2015). "Effects of siblings and birth order on income redistribution preferences: evidence based on Japanese General Social Survey." *Social Indicators Research*, 121: 589-606.