

Empirical research on economic inequality
Lecture notes on theories of justice (preliminary
version)

Maximilian Kasy

July 10, 2015

Contents

1	Considerations of justice and empirical research on inequality	3
1.1	Introduction	3
1.2	Individual welfare and social welfare	4
1.2.1	Approaches not covered by this framework	6
1.3	How to measure individual welfare?	7
1.3.1	Utilitarian welfare	7
1.3.2	Primary goods, capabilities	9
1.3.3	Opportunities	9
1.4	How to trade off welfare across people	10
1.4.1	Welfare weights	10
1.4.2	The veil of ignorance as an argument for equality	11
1.5	Why observable characteristics might matter	12
1.5.1	Equality of opportunity	13
1.5.2	Compensation for predetermined characteristics	13
1.5.3	Mechanisms	14
1.6	Summary	14
1.6.1	Recommendations for empirical research	14
1.6.2	Mapping notions of justice to questions in empirical research	17
1.6.3	Further readings	18

Chapter 1

Considerations of justice and empirical research on inequality

1.1 Introduction

This is a class about economic inequality. Such a class raises the obvious question: Why should we care about economic inequality? This question is closely related to the more general one, What makes a good society? Much ink has been spilled over the millenia on this question. This chapter cannot and will not attempt to provide a comprehensive overview of all that has been said on the subject. What we will do instead is discuss a general framework for the evaluation of “**social welfare**,” which measures how “good” a given society is, and some of the key questions that arise in this framework. Normative considerations arising in such a framework can provide guidance on what questions to ask about economic inequality, thus informing empirical research on the topic. Normative considerations are not the only reason to ask specific questions, of course; we might for instance also ask specific questions to better understand the mechanisms determining inequality, given a general interest in inequality.

The framework we will consider is individualist in that it first evaluates the **welfare of every individual** under consideration, and then **aggregates** to evaluate social welfare. Social welfare might be affected by policy. A policy change is considered desirable if it increases social welfare. Such a framework is consistent with many different views about social justice; some exceptions will be discussed below. Key questions arising in this framework include the following:

- **How should we evaluate individual welfare?**

In terms of achieved outcomes? In terms of resources at the disposition of individuals to achieve their objectives? In terms of options effectively

at their disposition? Relative to their own preferences or relative to some objective scale?

- **How should we aggregate**, that is, what is the relative importance assigned to the welfare of different people?
How much do we care about an additional dollar for a poor person versus a rich person? For a sick person versus a healthy person? Do we care more about inequality along dimensions such as race or gender than about other dimensions?
- What outcomes are considered **beyond the control** of an individual, and what outcomes are considered her own **responsibility**?
Should we be held responsible for outcomes predicted by our parents' economic status? Our skin color or gender? The nationality assigned to us at birth? Our level of education? Actions motivated by our religious beliefs?
- How should we deal with **characteristics that policy cannot** (or only with great difficulty) **affect**?
Should redistribution compensate for disabilities? For historical injustices? For education levels of adults?

This chapter obviously will not answer these questions. What we will do is discuss them using a formal framework. We will also relate them to various questions asked in empirical research on economic inequality. By doing so, we attempt to provide some coherence to the wide array of questions and methods considered in the literature and discussed in this book. Explicitly discussing these questions will also help you to be aware of some of the implicit normative assumptions involved in various empirical questions and methods.

At this point it should be emphasized again that normative considerations such as those discussed here are obviously not the only motivation for studying specific aspects of economic inequality. Conditional on being interested in economic inequality in general, we might also focus on specific dimensions to understand how inequality is produced and reproduced. For instance: What role does discrimination play? What role the resources provided by parents in early childhood? We might ask these questions, whether or not we believe that inequality along the dimensions of race, gender, or parental background is inherently worse than inequality along other dimensions.

1.2 Individual welfare and social welfare

As mentioned before, our framework is individualist. This implies that we have to consider a set of **individuals**

$$i = 1, \dots, n. \tag{1.1}$$

This already raises the first set of difficult questions: **Who is to be included** in this set of individuals? Many discussions implicitly assume we are considering a given “society.” But does that mean everybody of a certain citizenship, or everybody living in a certain territory? Why not all living human beings; should humans of another country count for nothing? And what about future generations? What about animals?

Given the set of individuals, we next need to decide how to measure their **welfare**. The goal is to assign a number

$$v_i \tag{1.2}$$

to each individual i , where v_i which measures how well they are doing. Depending on the setting, a stronger or weaker interpretation might be given to this number. An **ordinal** interpretation would only consider whether v_i is smaller or larger. A **cardinal** interpretation would care about the actual magnitude of v_i . An intermediate interpretation would consider the magnitude of changes of v_i . And the v s might or might not be **comparable across persons**. If they are, then it makes sense to say that v_i is bigger than v_j , that is, i is doing better than j .

How to measure individual welfare v_i again raises a whole set of difficult questions. A very minimalist notion would only consider the formal legal rights enjoyed by individuals. A broader notion might also take into account various resources that allow individuals to achieve their objectives, such as education, income, and health. A comprehensive notion of opportunities might aim to take into account all factors that influence individuals’ options, and evaluate the options effectively available to them. And we might finally consider the outcomes actually achieved by individuals, evaluated either by some common criteria, or by their individual preferences. Utilitarianism, the most common perspective in economics, evaluates individual welfare by the outcomes actually achieved as evaluated by individual preferences.

Given the set of individuals i , and given evaluations v_i of their welfare, we finally ask how well society as a whole is doing. Formally, we consider a “**social welfare** function,”

$$SWF = F(v_1, \dots, v_n). \tag{1.3}$$

The function F determines how much we care about different individuals. Note that everybody still has a “name” at this point; the function F tells us how much **weight** we assign to the **welfare** of i relative to the welfare of j , the welfare of Trevon relative to the welfare of Emily, the welfare of Sophia relative to the welfare of José. The function *might* treat different people similarly, and not care about names. In that case it would still tell us how much we care about an additional dollar for a poor person versus a rich person, for a sick person versus a healthy person, since these would affect the levels v_i . And the function F might not depend on names, but on some characteristics of individuals. The

function might for instance incorporate the belief that race, gender, or parental status should not determine individual welfare.

1.2.1 Approaches not covered by this framework

This individualist framework, where we evaluate social welfare SWF as a function of individuals' welfare v_i for $i = 1, \dots, n$, is very general. It is compatible with various perspectives, whether radical or conservative ones. It does impose some restrictions, however, excluding in particular both fascist and libertarian normative approaches. It also does not explicitly incorporate environmental concerns.

Our framework is not compatible with **fascist** approaches to the extent that these emphasize the greatness of a nation as their objective, no matter what the cost to the individuals involved. A slightly more benign variant of the fascist approach, sometimes called **perfectionist**, takes greatness in cultural production, science, etc. as its objective, again independently of the welfare of individuals. Such approaches have historically often been used to justify slavery or extreme inequality as necessary preconditions for the existence of a leisure class capable of cultural greatness.

Our framework is also not compatible with **libertarian** approaches in the tradition of John Locke. Such approaches consider outcomes to be just as long as they are the consequence of private property, contracts, and voluntary exchange on markets, *no matter what the consequences for individuals' welfare* are. Libertarian approaches are not the only justification for markets; market based organization of the economy might also be justified instrumentally, depending on its consequences for individuals. Such an instrumental perspective does not share the fundamentalism of the libertarian approach, however. In the instrumental perspective (which is consistent with our framework) market outcomes are evaluated based on how well different individuals are doing, and there might be all kinds of policy interventions that would improve social welfare if they are improving the welfare of at least some individuals, possibly at the cost of others.

Our framework, finally, does not explicit take into account **environmental** concerns. There are different ways to think about such concerns from a normative point of view. First, we might worry about degradation of the environment (man-made global warming, for instance) because of its impact on future generations. To the extent that individuals from future generations are included in our social welfare function, this is covered by our framework. Second, we might argue that the welfare of animals should carry some normative weight of its own. This would be covered by our framework if animals are included in the population considered. Third we might care about biodiversity, say, for its own sake rather than because of its impact on future generations or because of the welfare of individual animals. Such concerns are not directly covered by our framework. It is easily possible, of course, to consider objective functions F that depend both on the welfare of individuals and on other objectives, such as environmental concerns.

1.3 How to measure individual welfare?

In this section we discuss three alternative ways of thinking about individual welfare. The first one, which is the most common approach in economics, equates welfare to “utility,” and assumes that utility is what individuals maximize in their choices. The second one measures welfare in terms of the resources (“primary goods”) that individuals have in order to achieve their objectives, whatever these might be, or in terms of their “capabilities,” which take into account all constraints that individuals might face in pursuing their objectives. The third one attempts to give content to the idea of “opportunities.” The idea of this approach is to decompose individual outcomes into a component that individuals are deemed responsible for and a component that is determined by circumstances, called opportunities.

1.3.1 Utilitarian welfare

If you took a class in microeconomics at some point, you probably learned about utility and utility maximization. Evaluating individual welfare in terms of utility is by far the most common approach in economics. There are two main ingredients to this approach. First, a **choice set**

$$C_i \tag{1.4}$$

containing elements x . This choice set describes all options among which an individual can effectively choose, taking all constraints that she faces into account. Second, individual **preferences**, expressed in terms of a utility function

$$u_i(x). \tag{1.5}$$

The assumption is that individuals choose the element $x \in C_i$ which yields the greatest utility. The corresponding **level of utility** is taken to be a measure of individual welfare, so that

$$v_i = \max_{x \in C_i} u_i(x). \tag{1.6}$$

Note that utility plays two roles in this framework. Utility is (i) what individuals maximize when they have a choice, and (ii) what is deemed desirable for them from a social welfare point of view. It is generally assumed that preferences, that is, the utility functions u_i , are given independently of policy, while choice sets C_i might be affected by policy. Different policies might be compared in terms of the welfare v_i that individuals can achieve given these policies.

While this sounds straightforward enough, there is a key problem: **We don’t observe utility**. So how can we give empirical content to the idea of utility? Economists have come up with a neat trick to do so. Utility is not observable, but choices are. It is assumed that **choices are maximizing utility**. Leveraging this assumption, we can come up with an **equivalent** change in unearned income corresponding to any given change of the choice set. Such an equivalent change of income results in the same change of utility v_i as the given change of

the choice set. We will later discuss how this works in more detail. This trick allows us to calculate welfare effects on individuals in monetary terms. Once we have done so, what remains to decide is how much we care about additional dollars for different people in order to make statements about social welfare.

Things are of course not quite as simple in the real world. One key issue is that the utilitarian approach requires a neat separation of preferences and choice sets. Preferences are taken as the measuring stick by which to evaluate welfare, choice sets are affected by policy. But **preferences** of course don't exist in a pre-social vacuum, but are **affected by our social environment**. Important examples include professional aspirations affected by parental social class, and internalized gender role norms. This implies that it is hard to evaluate life chances at birth in terms of utility, since preferences are not given yet. How would one, for instance, evaluate a change in the educational system that affects the formation of professional aspirations?

Another issue, which has been emphasized by **behavioral** economists in recent years, is that individuals might have well defined predefined **preferences**, but **don't act accordingly** in all circumstances. There are various reasons why individuals might forget about important features of some choice because these features are not salient, why individuals might focus too much on the present relative to the future, or why individuals make bad choices because of stress and deprivation. The argument is that there is some "well considered" choice, corresponding to their "true preferences," which individuals would make if they would ponder carefully, but which they don't make in practice.

A third and very important issue is **comparability across individuals**. It is not obvious how we would evaluate whether Isabella or Ethan has a higher level of utility. But interpersonal comparisons are necessary to talk about inequality and to aggregate to measures of social welfare.

The historical evolution of the notion of utility

It is interesting to note how the notion of utility as a measure of welfare changed over time. **Originally**, as conceived by Bentham and his followers, utility was supposed to be a **cardinal** magnitude that was **comparable across people**, and in principle measurable. Utilitarianism called for maximizing the sum of individual utilities, thus answering both the questions of how to measure individual welfare and how to aggregate it to a notion of social welfare. This notion of social welfare had two main political implications. On the one hand, it incorporated a justification of **laissez-faire** economics: Whatever people do when left alone is what is best for them, by assumption. This is due to the double role of utility as a positive object (determining observable behavior) and a normative object (determining evaluations of welfare). On the other hand it carried radically **egalitarian** implications: Leaving aside the incentive effects of redistribution, if utility is comparable across people, and the value of additional income declines with the level of income, then social welfare is maximized by redistributing until everybody has the same income.

Later **revisions** of the notion of utility weakened the concept in a way the maintained the laissez-faire presumption, but did away with the egalitarian implications. Under the influence of positivist philosophy, utility was equated with the observable notion of **revealed preference**. Revealed preference says that the utility of option x is higher than the utility of option y if the given individual would choose x over y . This is the notion of utility taught in standard economics textbooks. This notion **maintains** the **laissez-faire** presumption that whatever people do when left alone is what is best for them. It however reduces utility from a cardinal notion that is comparable across people to an ordinal notion that is not comparable. Neither the notion that the marginal utility of income is declining, nor comparisons of welfare across people make sense from the perspective of revealed preference. The **egalitarian implications** thus **disappear**. A notable exception to this general evolution is modern public finance, and in particular the theory of income taxation, which largely is based on the assumption that utility can be compared across people.

1.3.2 Primary goods, capabilities

The utilitarian approach allows to evaluate what is better or worse for a given individual, based on what they would choose if they could. It implies, in particular, that more options (a larger choice set C_i) are generally better. It does not, however, allow to compare utility across people, at least based on revealed preference. Being able to do such a comparison is key for talking about inequality from a normative point of view.

What we can observe and **compare across people** are the **resources** that they have at their disposition, and that affect the choices effectively at their disposition. More resources generally means more options and thus higher utility. Considering these resources is one of the key ideas in Rawls (1973), who calls them **primary goods**. He argues that health, civil and political rights, income and wealth, and social bases of self-respect should be among the primary goods by which individual welfare is to be measured. This approach allows to say who is doing better or worse in terms of resources, and thus whether equality is achieved or not.

A similar notion, following up on Rawls, was developed by Sen (1995). He argues that the key measure of welfare are “**capabilities** to function.” Rather than being based on a fixed list of resources or primary goods, capabilities are comprehensive measures of the options effectively at the disposition of individuals, taking into account all legal, social, economic, and cultural constraints. These capabilities are quite similar to the sets C_i introduced above in our discussion of utilitarian welfare.

1.3.3 Opportunities

A key question that tends to come up in discussions about inequality is the question to what extent differences in outcomes should be considered the **responsibility** of the individuals involved, or the consequence of **social constraints**.

Many people would subscribe to the idea of “equality of opportunity,” where policy should not necessarily equalize outcomes, but should “level the playing field.”

It is surprisingly difficult to formalize this idea; an attempt to do so can be found in Roemer and Trannoy (2015). The key issue is similar to the difficulty of separating preferences from choice sets. In the equality of opportunity context this is framed as separating **effort** from circumstances. Inequality conditional on **circumstances**, determined by effort, is considered legitimate, while inequality determined by circumstances is considered bad. Circumstances are also called **morally arbitrary**; they should not affect outcomes. To put things into our framework, we could set welfare v_i to be life outcomes as predictable by circumstances. Roemer and Trannoy (2015) argue that the separation between the two categories, effort and circumstances, should be based on social convention. Opportunities are quite similar to capabilities. The latter are attempting to give a principled description of the options effectively available to someone subject to *all constraints*, however, while opportunities only consider a more *restrictive list of constraints* based on social convention, and consider the rest to be an individual’s responsibility.

To illustrate the difficulty in separating responsibility from circumstances, suppose someone is born into a religious community with norms that imply a life in poverty. Should this poverty be considered a matter of choice that is to be respected but is not deserving of redistribution, or a consequence of unequal opportunities at birth? Or, somewhat more subtly: Suppose social norms are such that girls on average end up pursuing educational tracks that lead to lower paying professions and lower labor market participation. Is this a matter of opportunities or choices?

However we answer these questions, it appears evident that certain things are beyond our control: our race, the social background of our parents, the nationality that we are born with, or the gender assigned to us. One way to formalize, if not equality of opportunity itself then at least a necessary condition, is to say that such characteristics determined at the time of our birth should not be predictive of our outcomes later in life. Arguably, the more variables we consider when defining circumstances, the better we are able to predict life outcomes at birth. When we consider a comprehensive set of pre-determined variables, there might in fact not be much of a difference left between equality of opportunity and equality of outcomes.

1.4 How to trade off welfare across people

1.4.1 Welfare weights

In the previous section we discussed various perspectives on how to measure individual welfare. Alternative options include the realized utility of an individual, the primary goods at her disposition, capabilities, or opportunities. Whichever way individual welfare is measured, in order to make normative statements we

must aggregate individual welfare to social welfare, using some function F which takes care of this aggregation.

Suppose we are interested in small changes of some policy, which results in corresponding small changes dv_i of welfare v_i for each individual. It is then useful to think about aggregation in terms of **welfare weights** ω_i : Recall that social welfare is equal to $SWF = F(v_1, \dots, v_n)$. Assume that F is differentiable and let

$$\omega_i := \frac{\partial}{\partial v_i} F(v_1, \dots, v_n). \quad (1.7)$$

We can calculate the change $dSWF$ of social welfare by taking derivatives,

$$dSWF = \sum_i \omega_i \cdot dv_i. \quad (1.8)$$

where dv_i is the effect of the policy change on the welfare of individual i and ω_i measures the weight attached to the welfare of individual i when calculating social welfare.

In a utilitarian context, dv_i is often expressed in monetary terms (as an amount of US\$, for instance). dv_i is then called equivalent variation. If dv_i is measured in monetary terms, then the ratio ω_i/ω_j measures how much we care about an additional US\$ for person i relative to an additional US\$ for person j . An egalitarian position assigns a large weight ω on those who are worst off in a society. More generally, we can think of the ratio between welfare weights for poor versus rich people as measuring how egalitarian some aggregation is. Utilitarianism, as originally conceived, assumes that the marginal utility of income is declining with income, and thus puts a higher weight on those worse off.

In contrast to such an aggregation, many models used by economists, are just **summing up US\$ across people**, thus assigning the same weight to additional income for everyone. An additional US\$ for a billionaire is considered equivalent to an additional US\$ for a starving person. Examples of such aggregation are conventional calculations of consumer surplus, and of dead weight loss. Such an aggregation is very **anti-egalitarian**, and has no theoretical justification; you should be skeptical when you see a model which does this.

We can also think about **welfare weights “in reverse:”** Given the policy choices actually made in a society, there is a corresponding set of welfare weights which justifies these policy choices. Such weights can be thought of as measuring **effective social power** – whose interests are represented by the powers that be.

1.4.2 The veil of ignorance as an argument for equality

The idea the human beings should be treated symmetrically is an old one – “Do unto others as you would have them do unto you.” This idea of **symmetry** (or **impartiality**) is closely related to the idea of **equality**. If everybody’s welfare counts the same, so that the subscript i of v_i does not matter, then it is natural that we should aim to improve the welfare of those worse off, possibly with a

cost to those doing well. As argued by Sen (1995), almost all modern theories of justice demand equality of some sort or other; they do differ however in terms of their answer to the question Equality of what? Put differently, there is more disagreement about the definition of v_i than about the choice of F .

But is there a deeper **justification for egalitarianism** beyond the fact that it is quite hard to justify treating people unequally along the dimension “that matters?” Rawls (1973) famously provided such a justification with his thought experiment of the **veil of ignorance**. This thought experiment provides a twist on the old idea of a social contract, reducing the evaluation of social welfare SWF to a **decision problem under uncertainty**.

The thought experiment goes as follows. Imagine you knew nothing about yourself, just that you are some person in a given society. The next morning you could wake up as any person in that society, you could be any $i = 1, \dots, n$. Suppose you have to choose between various alternative social arrangements and policies. Which one would you pick? This point of view, where you don’t know who you are, is also called the **original position**.

Rawls argues that, when confronted with this problem of choice under fundamental uncertainty, then we should uphold two principles for a desirable society. First, everyone should be free to choose their own objectives in life. Second, we should make sure that those who are worst off are as well off as possible. This second principle is also called **maximin** (maximizing, by choice of policy, the minimum of welfare, across individuals). Formally we could write

$$F(v_1, \dots, v_n) = \min_i v_i. \quad (1.9)$$

This choice of F implies that a society is only as good as the fate of those who are worst off in it. Such a formulation requires that we are able to compare v_i across individuals to determine who is in fact worst off. The reason for choosing such an F is that, when in the original position, we are faced with the possibility that we might end up as one of those who are worst off. Picking social arrangements which maximize this minimum is a way of insuring ourselves against this possibility. In terms of welfare weights, the maximin specification assigns a much larger weight to those who are worst off relative to everyone else.

1.5 Why observable characteristics might matter

There is some research on economic inequality which is concerned with inequality overall (whether of incomes, earnings, wages, wealth, consumption, or some other magnitude). A lot of research, however, focuses on inequality along some specific dimension as opposed to inequality overall: Inequality by race or ethnic group, by gender, by parental background, by educational attainment, by nationality, by location of residence, etc. There are various reasons for focusing on these dimensions, and these reasons are not mutually exclusive. In this section we review some of these reasons.

1.5.1 Equality of opportunity

A first reason for caring about inequality along some observable dimension is that this dimension might be considered as reflecting circumstances beyond an individuals' control. We discussed this possibility when considering "opportunities" to be a measure of individual welfare. To the extent that we believe that equality of opportunity along such a dimension is desirable, it is normatively relevant to study inequality along such a dimension, its historical change, or how it is affected by policy choices.

Empirically, we might in particular wish to study inequality by race or ethnic group, by gender, by parental background etc. because such inequality is considered particularly undesirable.

1.5.2 Compensation for predetermined characteristics

A second, related, reason for caring about inequality along some observable dimension is the following. Suppose that the characteristic under consideration cannot be directly affected by policy, but we might wish to compensate for the effect on welfare that this characteristic has. Obvious example are disabilities and (innate) abilities.

Are there reasons to compensate? Dworkin 2002 has proposed the following thought-experiment which is a variant of Rawls' veil of ignorance. Suppose individuals knew some of their characteristics, but not all of them. They don't know their abilities and handicaps, in particular. They have to negotiate and contract to decide on a society and sets of policies to be implemented. The argument is that they would come up with **insurance schemes** to compensate for all the characteristics hidden behind the "**thin veil of ignorance**." According to Dworkin, the redistributive institutions of the welfare state are a way to effectively implement such insurance schemes.

The role of things hidden behind the thin veil of ignorance in this thought-experiment is similar to the role of circumstances in discussions of equality of opportunity. What is hidden behind the veil is morally arbitrary. Note, however, that the insurance argument for compensation for predetermined characteristics is also consistent with a purely utilitarian conception of individual welfare (and with other notions of welfare). Markets are by definition unable to provide the kind of insurance Dworkin has in mind, because it is insurance for pre-determined characteristics.

Can pre-existing differences be compensated? Does the fact that inequality is largely pre-determined at birth not imply that policy can do little about it? Goldberger (1979) has proposed the following famous example regarding genetic heritability. Suppose we are interested in eyesight. It is well possible that eyesight is 99% genetically determined. Now consider a policy which distributes glasses for free to everyone. It is again quite conceivable that such a policy completely equalizes effective eyesight, at low cost, increasing everyone's welfare. What this example teaches us is that what matters is the **interaction of predetermined characteristics with social institutions in generating inequality**.

1.5.3 Mechanisms

We might, finally, want to study inequality along particular dimensions because this allows us to gain a better understanding of the mechanisms by which economic inequality is produced and reproduced, and to understand how policies might affect inequality. Let us consider some examples

There is a large literature in labor economics which studies the returns to education (that is, the inequality between workers of different levels of education), how returns have evolved over time, and how they are affected by the supply of workers of various levels of education. One might care about returns to education not because they seem inherently more unjust than inequality along other dimensions, but because they suggest obvious ways in which policy might affect inequality. Expanding public higher education, for instance, might affect inequality by (i) allowing additional people to benefit from the returns to education, (ii) reducing these returns through the increase in the relative supply of highly educated workers, and (iii) through redistribution from taxpayers to those attending universities and colleges.

Another literature in economics studies racial discrimination in the labor market. In addition to normative considerations suggesting that racial inequality is particularly unjust, studying the extent and mechanisms of discrimination in hiring allows us to better understand the extent to which inequality in general and racial inequality in particular might be reduced by anti-discrimination legislation or other measures.

1.6 Summary

1.6.1 Recommendations for empirical research

In this chapter, we have provided a brief overview of some of the issues and debates in the theory of distributive justice. Without wishing or being able to resolve any of these debates, we believe there are still some general recommendations that can be derived for good empirical research which is conscious of these debates.

These general recommendations aim to (i) produce results which are relevant for readers holding various views on distributive justice, and (ii) make explicit normative assumptions often kept implicit. Both (i) and (ii) are important in light of the fact that many traditional approaches in economics are based on questionable normative assumptions, such as aggregation by summing up dollars across people, or equating income and welfare.

Aim to report disaggregated results

We described the evaluation of social welfare as involving two steps. First we need to evaluate individual welfare v_i for all individuals. Then we need to aggregate using the function F to get $SWF = F(v_1, \dots, v_n)$. People disagree on the appropriate choice of F ; they have different opinions about how much the

welfare of different individuals i should matter. If we, however, know the **distribution of the** v_i , joint with other relevant characteristics of individuals, then we can **evaluate social welfare no matter what** F we choose.

Empirical papers that report disaggregated results on the distribution of v_i , or on the effect of policies on this distribution, thus allow the reader to evaluate based on their own normative judgment regarding F . Reporting disaggregated results makes explicit the distributional dimension involved in virtually all policy decisions and historical developments. In general there are always **winners and losers**.

Reporting disaggregated results also avoids the anti-egalitarian assumptions involved in many traditional models in economics. As discussed before, in many traditional models, social welfare is equated to a summing up of utilities measured in dollars. This is arguably the least egalitarian welfare function conceivable, where an additional dollar for a billionaire is valued as much as an additional dollar for a starving person.

How can results be reported in a disaggregated way? One good option are **quantiles**. The τ^{th} quantile of v is the value Q^τ such that a share of τ of the population is below Q^τ ,

$$P(v_i \leq Q^\tau) = \tau.$$

Quantiles for low values of τ tell us how the poor are doing, quantiles for large values of τ tell us how the well-off are doing. For $\tau = 0.5$, we get the median of the distribution of v . In addition to reporting quantiles, it is also good practice to report **results separately** for various **demographic groups**, defined for instance by race, gender, country of origin, location of residence, etc.

Be aware that your variable probably does not measure welfare

As we discussed, the most important difference between various perspectives on social welfare and inequality is how individual welfare v_i is measured. Prominent proposals include utility, primary goods, capabilities, and opportunities. It is not always easy to map these concepts to observable magnitudes. Empirical research on economic inequality often focus on more easily measured variables such as income, earnings, wages, wealth, or consumption.

Studying these variables is of course relevant and interesting. It is important, however, to be conscious that they are not direct measures of welfare. Consider for instance earnings, maybe the most studied of these variables. There are many policies which might increase the earnings of low income workers, and studies which find that they do usually suggest these policies to be desirable. Often, however, earnings are increased by making the life of the unemployed more miserable, whether by reducing transfer payments or through the increase of bureaucratic hassles and ordeals which they have to undergo. There is arguably no notion of individual welfare for which such a policy change is increasing welfare for the unemployed.

There are many related examples which are more subtle. Earnings subsidies might for instance increase hours worked, thereby increasing earnings without necessarily increasing utility (or other notions of welfare) to the same extent.

Another example is the increased labor force participation of women which is coupled with increased marketization of care work. Increased earnings go along with increased expenditures for care. Such an increase in labor market participation might of course be considered desirable from a perspective of gender equality, shifting intra-household bargaining positions. It also affects the situation of those (often migrant women) who perform the marketized care work. But the increase in earnings is not by itself an increase in welfare.

Be explicit why you study a specific dimension of inequality

As we discussed above, there are several reasons why one might study inequality along a particular dimension. A dimension might be relevant because we think that, from an equality of opportunity perspective, this dimension should not matter for outcomes. It might also be relevant because we think that public institutions could and should compensate for the disadvantages implied by some predetermined characteristics, and because we want to understand the interaction of these characteristics with social institutions in producing outcomes. We might finally care about some characteristics not because we think that inequality along these characteristics is inherently more unjust than inequality along other dimensions, but because this allows to gain a better understanding of some of the determinants of inequality overall.

It is a good idea to explicitly discuss **why** one studies a specific dimension of inequality. This entails in particular an acknowledgment of the **limitations of studying only this dimension**. Some examples again help to illustrate.

There is a sizable literature that studies discrimination in hiring, along dimensions such as race or gender. This literature focuses on differences in the probability of getting invited for a job interview of (hypothetical) applicants who are identical except for their race, say. This is of course interesting. The danger is to equate the absence of discrimination in hiring with equality across groups, and to equate equality across groups with equality overall.

There are many channels through which group inequalities might be created. These include early childhood influences, different neighborhoods of growing up, different access to and quality of primary, middle, and high school education, the creation of aspirations, different access to and treatment in higher education, different chances of being hired when applying for a job, different wages conditional on being hired, different chances of being promoted or fired in a given job, differential treatment by customers or clients, etc. The channel of *hiring* might in turn be decomposed into several components. What is the chance of being invited to an interview, and what is the chance of being hired given an interview? How does the chance of being invited to an interview depend on the neighborhood of residence, the high school attended, or the (perceived) race and gender of an applicant? Differences in the probability of getting an interview are just **one of the many channels** creating inequality between groups.

Another large literature studies intergenerational mobility, that is, the correlation between incomes (or other outcomes) of parents and children. This is interesting and tells us a lot about the pre-determination of life outcomes at

birth. The pitfall is, once again, to equate the absence of such a correlation with equality of opportunities. There are many factors and channels through which life outcomes are affected and which are beyond an individual's control which are independent of parental income.

1.6.2 Mapping notions of justice to questions in empirical research

We conclude by providing a list of topics which we will discuss in our review of econometric methods for the study of economic inequality. For each of these topics, we mention key concepts from theories of social justice that relate to this topic.

1.
 - Topic: The long run evolution of inequality as measured by top income shares
 - Concepts of social justice:
Inequality of outcomes.
The role of society in determining (the distribution of) outcomes, as demonstrated by historical variability.
2.
 - Topic: The long run evolution of gender inequality
 - Concepts of social justice:
Gender as a morally arbitrary category.
Interaction of labor markets and care institutions with gender to produce inequality of outcomes.
3.
 - Topic: Racial discrimination
 - Concepts of social justice:
Race as a morally arbitrary category.
Labor market discrimination as one mechanism generating racial inequality.
4.
 - Topic: The effect of de-unionization on inequality
 - Concepts of social justice:
Unionization as mechanism affecting inequality.
5.
 - Topic: Labor demand and labor supply, technical change, immigration
 - Concepts of social justice:
Labor demand and labor supply as mechanism affecting inequality.
6.
 - Topic: Intergenerational mobility
 - Concepts of social justice:
Parental background as circumstance determining opportunities.
Policies compensating for parental background.
Intergenerational transmission as mechanism producing and reproducing inequality of outcomes.

7.
 - Topic: The welfare impact of changing prices and wages
 - Concepts of social justice:
 - Utility as a measure of individual welfare.
 - Equivalent variation as change in utility expressed in monetary terms.
8.
 - Topic: Redistributive taxation
 - Concepts of social justice:
 - Utility, equivalent variation.
 - Welfare weights depending on income.
 - Maximization of social welfare.
9.
 - Topic: International inequality
 - Concepts of social justice:
 - Whose welfare is to be included in the social welfare function?
 - Nationality as a circumstance determining opportunities.

1.6.3 Further readings

Readers interested in the normative questions raised here might want to study various other references. The point of departure for much of the modern philosophical literature on social justice is the following:

Rawls, J. (1973). *A theory of justice*. Harvard University Press, Cambridge.

Rawls argues for equality of the distribution of so-called “primary goods” based on a construction called “veil of ignorance;” we have reviewed this argument above.

A discussion of various perspectives on justice and inequality is provided by

Sen, A. (1995). *Inequality reexamined*. Oxford University Press, Oxford.

Sen argues that most modern theories of justice and inequality are egalitarian, including conservative theories. These various theories differ, however, in their answer to the question Equality of what? Sen proposes to answer this question by “capabilities to function.” What counts, according to Sen, is the set of options among which individuals can effectively choose, taking into account all consequences, and all legal, economic, social, and cultural constraints that they face.

Another ambitious review of the literature on justice and the evaluation of social welfare can be found in

Roemer, J. E. (1998). *Theories of distributive justice*. Harvard University Press, Cambridge.

Roemer discusses two related literatures, social choice (a subfield of economic theory), which tends to be very formal and axiomatic, and political philosophy. Roemer attempts to bring insights from both literatures to bear on each other.

Bibliography

- Dworkin, R. (2002). *Sovereign virtue: the theory and practice of equality*. Harvard University Press.
- Goldberger, A. S. (1979). Heritability. *Economica*, 46(184):327–347.
- Rawls, J. (1973). *A theory of justice*. Harvard University Press, Cambridge.
- Roemer, J. E. (1998). *Theories of distributive justice*. Harvard University Press, Cambridge.
- Roemer, J. E. and Trannoy, A. (2015). Equality of opportunity. In Atkinson, A. B. and Bourguignon, F., editors, *Handbook of Income Distribution*, volume 2 of *Handbook of Income Distribution*, chapter 4, pages 217 – 300. Elsevier.
- Sen, A. (1995). *Inequality reexamined*. Oxford University Press, Oxford.