

Distributive Equality

DAVID MCCARTHY

University of Hong Kong

mccarthy@hku.hk

Egalitarians think that equality in the distribution of goods somehow matters. But what exactly is egalitarianism? This article argues for a characterization based on novel principles essentially involving risk. The characterization is then used to resolve disputed questions about egalitarianism. These include: the way egalitarianism is concerned with patterns, in particular its relationship to strong separability; the relationship between egalitarianism and other distributive views, such as concerns with fairness and with giving priority to the worse off; and the relationship between egalitarianism and evaluative measurement. But egalitarianism is subject to a particularly severe form of the levelling-down objection, and is claimed to be false.

Introduction

Egalitarians think that equality in the distribution of goods somehow matters. Much has been written about particular forms of egalitarianism. For example, it has been asked whether equality becomes more important as people become better off, or whether it becomes less important. This article addresses a different question: What exactly is egalitarianism?

The article breaks naturally into two parts. The first part, in sections 1 to 9, offers an account of the core of egalitarianism, then embeds this core into some background ideas to form what it calls theory E. Ignoring certain simplifying assumptions, the main claim of the article is that egalitarianism is theory E. The second part, in sections 10 to 21, then explores what that claim teaches us about various egalitarian debates. Although the main goal is to understand egalitarianism, the article will also try to assess its truth.

In more detail, the first part starts with a natural approach to understanding egalitarianism. This approach assumes that we can measure well-being quantitatively, and does not say anything about risk. It then criticizes this approach to motivate trying to understand egalitarianism without talking quantitatively about well-being, and instead by talking about risk. After rehearsing some background to

do with expected utility theory, it develops an account of the core of egalitarianism, eventually adding background ideas to form theory E. It then shows how natural egalitarian theories can be formed by strengthening theory E in various ways, and claims that its basic approach to egalitarianism is stable under revisions of detail.

The second part looks at specific questions about egalitarianism, partly to try to answer those questions, and partly to test the claim that egalitarianism is theory E. These include: How does theory E relate to more traditional characterizations of egalitarianism; and how does egalitarianism, understood as theory E, relate to other distributive views? But the most discussed topic is the way egalitarianism is concerned with patterns between people. Contrary to a widespread view, it is argued that egalitarianism is compatible with a notion known as strong separability. Instead, egalitarianism should partly be characterized as denying a little-discussed form of separability I call weak separability across individual lotteries. But this denial opens egalitarianism up to something I call the strong levelling-down objection. The article ends by looking at five ways egalitarians might respond, finding them wanting, and concluding that egalitarianism is false.

1. Luck and distribution

The goal of the first part of the article is to offer an account of egalitarianism as a view about when one distribution of goods is better than another. While there may be ideas about topics aside from distribution which should be seen as egalitarian, I am not concerned with those. But one requires further comment.

Suppose we can distinguish between luck which is the result of a deliberately chosen gamble which an agent might reasonably have declined, and any other form of luck. For example, Dworkin, (2000, p. 73) calls the former ‘option luck’ and the latter ‘brute luck’. Luck egalitarianism is often presented as the view that a distributive ideal is the elimination of the effects of brute luck. For example, Cohen (1989, p. 908) goes so far as to regard that elimination as ‘the primary egalitarian impulse’. Given that how lotteries resolve is a paradigm form of luck, it would seem that there is already a sophisticated literature analysing egalitarianism in terms of risk.

However, despite the way luck egalitarianism is often presented, it is highly implausible that equality can somehow be usefully analysed in terms of the elimination of brute luck. Equality can easily arise as a

matter of extreme luck (consider a lottery which offers only a tiny probability of an equal outcome), and it seems completely indeterminate what sort of situation would arise in the absence of any luck (see Hurley 2001). As a corollary, there is little prospect for arguing that equality is a good thing on the grounds that brute luck is somehow or other a bad thing.

Moreover, despite sometimes incautious advertising, luck egalitarians are typically not concerned with explicating or defending the claim that equality is valuable. Rather, they are concerned with what they take to be the obverse of brute luck, namely responsibility. Thus as Arneson (2001) points out, luck egalitarians typically just assume that equality is valuable, and are concerned with the extent to which departures from the ideal of equality could be justified by claims about responsibility. For example, if someone is in some sense responsible for shrinking the social pie, is her claim on an equal share of the remainder correspondingly diminished?

My concern is not with responsibility, but rather with what it means to say that equality is a distributive ideal in the first place, and whether it really is a distributive ideal. So I will be saying relatively little about luck egalitarianism, and, to avoid distractions, I will assume no one has any responsibility for the situations I will be discussing. They can all be seen as involving manna from the sky.

2. The fundamental egalitarian idea

As a distributive theory, I take the following to be the fundamental egalitarian idea (in a preliminary version):

The fundamental egalitarian idea, preliminary version:

Two things matter: increasing well-being, and decreasing inequality

My goal is to shape this suggestive but vague picture of egalitarianism into something more precise. I assume a fixed population of individuals $1, \dots, n$ throughout.

Why not simplify further and say that egalitarians only care about decreasing inequality? This leaves out too much. If egalitarians did not value well-being, it is hard to see why they would care about how it is distributed. Also, egalitarians will surely think that everyone equally well off is better than everyone equally badly off.

I am using 'well-being' as a cover-all term, intended to be neutral between competing accounts of what egalitarians should care about the distribution of. Familiar accounts have to do with happiness,

desire satisfaction, objective goods, and so on. I will also use the term ‘individual goodness’ more or less synonymously. It has the advantage of sounding more neutral, but it is awkward and I will tend to speak of well-being, or of one person being better off than another, when being informal.

This section now outlines one natural attempt to shape the fundamental egalitarian idea into a more precise characterization of egalitarianism. The next section will then criticize this attempt to motivate a different approach.

The following is a natural sharpening of its preliminary version.

The fundamental egalitarian idea:

- (a) If two situations are equally good in terms of well-being, and one is better than the other in terms of equality, it is better
- (b) If two situations are equally good both in terms of well-being and in terms of equality, they are equally good
- (c) If one situation is better than another in terms of well-being, and at least as good in terms of equality, it is better

When I say, for example, ‘better in terms of equality’, I mean something like ‘better from the point of view of a sole concern with equality’.

While the preliminary version seems obvious, if vague, this version may not. Nevertheless, it grows quite naturally out of the preliminary version. Clause (b) essentially says that nothing matters apart from well-being and equality. I do not deny that there could be a hybrid theory which cared about well-being, equality, and some further value. But here I am going to restrict attention to pure versions of egalitarianism which are only concerned with those two values. Of course, (b) does not imply that egalitarians do care about well-being or equality. For example, a theory which in effect does not care about anything, and regards all situations as equally good, could accept (b). Thus (a) has the effect of saying that egalitarianism values equality, while (c) says that it values well-being.

It is important to forestall an immediate objection. What I am calling the fundamental egalitarian idea only addresses situations in which the two egalitarian concerns, well-being and equality, never disagree. But it may be objected that what egalitarians are really concerned with is when the two concerns point in opposite directions: when well-being favours one situation and equality favours another.

This is the kind of conflict which lies at the heart of egalitarianism, so a story about egalitarianism has to address it directly, and help us to understand how egalitarianism thinks about trade-offs between well-being and equality.

In response, suppose a concern with well-being favours a situation s_1 while a concern with equality favours s_2 . Egalitarianism is by itself not going to tell us which situation is on balance better. One does not fail to be an egalitarian merely by giving less weight to equality than well-being, and some such view will favour s_1 on balance; nor does one fail to be an egalitarian merely by giving more weight to equality than well-being, and some such view will favour s_2 .

However, this article aims to give an account of what egalitarianism amounts to, not to take sides in disputes between different versions of egalitarianism. So albeit somewhat vaguely at this stage, what I am calling the fundamental egalitarian idea insists only that egalitarians accept the seemingly bland ideas involved in the cases where there is no conflict between well-being and equality. To say how conflict cases should be resolved would stray beyond neutrality.

To develop the fundamental egalitarian idea we face a problem. There is dispute among egalitarians about when one situation is better in terms of equality than another. But take one of the clauses of the fundamental egalitarian idea, such as clause (a). Egalitarians will agree that if two situations are equally good in terms of well-being, and one is better than the other in terms of equality, it is better. But they will not in general agree about when one situation is better than another in terms of equality. So they will not in general agree about what (a) really amounts to. Thus if something along the lines of the fundamental egalitarian idea is to be offered as a characterization of egalitarianism, we face a difficulty. If we try to develop it by offering a concrete account of when one situation is better than another in terms of equality, we are bound to end up excluding some genuine forms of egalitarianism. On the other hand, without a concrete account, we risk the content of the eventual characterization being too indeterminate to be interesting or informative.

A natural strategy for dealing with this problem focuses on what egalitarians do agree about. For although some judgements about betterness in terms of equality are controversial among egalitarians, many are not. Thus it is uncontroversial that, for example, a situation of perfect equality is better in terms of equality than a situation of inequality. Thus say that the *minimal betterness in terms of equality relation* holds between two histories h_1 and h_2 (or alternatively, h_1 is at

least as good in terms of minimal equality as h_2) just in case h_1 is at least as good in terms of equality as h_2 according to any genuine form of egalitarianism. By a *history* I mean an entire world-history, or possible world, past, present, and future.

Suppose we can provide a good account of the content of the minimal betterness in terms of equality relation. Then by considering the ways in which minimal betterness in terms of equality interacts with betterness in terms of well-being, we should be able to arrive at a minimal but determinate set of claims about overall betterness which every form of egalitarianism is committed to. The idea, then, is that accounts of overall betterness will be egalitarian just in case they accept this minimal set of claims, leaving room for them to disagree about other claims. They can disagree on what else to say about betterness in terms of equality beyond the minimal betterness in terms of equality relation, and they can disagree on the relative weights to give to equality and to well-being. Of course, this is all very rough, and only indicates a strategy with no guarantee of success. But guided by the fundamental egalitarian idea, we can now start to make it more concrete.

To simplify, I will focus on understandings of equality in terms of how well people's entire lives go. So I will ignore issues about equality at, say particular times discussed by McKerlie (1989), although I believe my approach could handle them. Say that an *individual goodness measure* is a quantitative measure of how good histories are for particular individuals. Assuming that individual goodness measures are well-enough defined for talk of sums of individual goodness to be meaningful, the following seems natural.

The risk-free interpretation of the fundamental egalitarian idea:

For any histories h_1 and h_2 , the following hold:

- (a) If h_1 contains the same sum of individual goodness as h_2 , and is better in terms of minimal equality, then h_1 is better than h_2
- (b) If h_1 contains the same sum of individual goodness as h_2 , and is exactly as good in terms of minimal equality, then h_1 and h_2 are equally good
- (c) If h_1 contains a greater sum of individual goodness than h_2 , and is at least as good in terms of minimal equality, then h_1 is better than h_2

For example, letting the numbers measure individual goodness, the risk-free interpretation will imply that according to egalitarianism:

$[1, 1]$ is better than $[2, 0]$; $[1, 0]$ and $[0, 1]$ are equally good; and $[1, 1]$ is better than $[0, 0]$.

The risk-free interpretation seems to capture the fundamental egalitarian idea in a natural way. It is widely held that utilitarianism cares about well-being but is indifferent to matters of distribution. Thus from the point of view of a sole concern with well-being, one history is at least as good as another just in case it contains at least as great a sum of individual goodness. The risk-free interpretation just slots this idea into the fundamental egalitarian idea, and takes situations to be histories.

The risk-free interpretation falls short of offering a full characterization of egalitarianism in three obvious ways. First, it does not offer an account of when one history is at least as good in terms of minimal equality as another. Second, a full characterization of egalitarianism will no doubt draw on various additional background assumptions, such as the idea that betterness is transitive. Third, it does not tell us how egalitarianism applies in conditions of risk, and a full characterization will eventually have to deal with that. But despite the fact that there is still work to be done, the risk-free interpretation does seem to be a strong candidate for expressing the core of egalitarianism.

3. Criticisms of the risk-free interpretation

In assuming that sums of individual goodness can be compared, the risk-free interpretation takes it for granted that it makes sense to talk about units of individual goodness. For such talk to be well-defined, individual goodness measures have to be unique up to so-called *positive affine transformation*, that is, multiplication by a positive number and addition of some arbitrary number. But the assumption that they are unique in that way runs into trouble because of a simple but important mathematical fact. We will later need this fact in several contexts.

A binary relation R on a set X is *transitive* if for all $x, y, z \in X$, xRy and yRz implies xRz , and *complete* if for all $x, y \in X$, xRy or yRx . An *ordering* is a binary relation which is transitive and complete. A real-valued function f represents an ordering \succsim if for any x and y in the domain of the ordering, $x \succsim y$ iff $f(x) \geq f(y)$. Consider the following, which shows that as representations of orderings, functions are only unique up to increasing transformation.¹

¹ If $f : X \rightarrow Y$ and $g : Y \rightarrow Z$ are functions, the function $g \circ f : X \rightarrow Z$ is defined by $g \circ f(x) := g(f(x))$. The expression ' $A := B$ ' means that A is defined as B . A function

Lemma 1

For any ordering \succsim and strictly increasing function h : a function f represents $\succsim \Leftrightarrow h \circ f$ represents \succsim

Say that the *risk-free extended individual betterness relation* holds between two individual/history pairs (i, h_1) and (j, h_2) if: h_1 is at least as good for i as h_2 is for j . The only obvious constraint a function g on individual/history pairs has to satisfy to be an individual goodness measure is that it must represent the risk-free extended individual betterness relation. But by Lemma 1, this only forces individual goodness measures to be unique up to increasing transformation. But there are vastly more increasing transformations than positive affine transformations, so the risk-free interpretation seems to suffer from pre-supposition failure.

There is a lot to say about this problem, but for now I will make two points. First, Broome (2004) argues that it is just vague how to measure how good things are for people. If we are going to use individual goodness measures, we need to adopt some sort of arbitrary definition. Suppose we do that. It then follows that for all we know, any principle, like the risk-free interpretation of the fundamental egalitarian idea, whose content is sensitive to the way individual goodness is being measured, will rest on arbitrary foundations.

For example, assume a population of two people with equality in h_1 , and the first person better off than the second in h_2 . We could arbitrarily fix an individual goodness measure which would give these histories individual goodness profiles of $[1, 1]$ and $[2, 0]$ respectively. Relative to that choice the risk-free interpretation implies that according to egalitarianism, h_1 is better than h_2 . But we could also arbitrarily fix an individual goodness measure which gave them individual goodness profiles of $[1, 1]$ and $[3, 0]$. The risk-free interpretation no longer implies that according to egalitarianism, h_1 is better than h_2 . But this arbitrariness should not satisfy anyone who thinks that egalitarianism is a natural and ethically fundamental distributive theory. Now perhaps Broome is wrong about individual goodness measures. But at the very least this shows that it is unsafe to follow the risk-free interpretation in regarding individual goodness measures as given, and then using them to develop an account of egalitarianism.

$f : \mathbb{R} \rightarrow \mathbb{R}$ is strictly increasing if for any $x, y \in \mathbb{R}$, $x > y \Rightarrow f(x) > f(y)$. Thus strictly increasing functions preserve the order of the real numbers, which is basically what Lemma 1 says. More precisely, proof of the lemma follows from the fact that if f and h are as per Lemma 1, $f(x) > f(y) \Rightarrow h(f(x)) > h(f(y))$.

Second, the risk-free interpretation was motivated by the widely held idea that utilitarianism is indifferent to equality. But with many others (e.g. Broome 1991, Hammond 1991), I take the premisses of the best defence of utilitarianism to be supplied by a version of the famous aggregation theorem of Harsanyi 1955. But the premisses of this theorem make no mention of individual goodness measures, so they could be accepted by someone who thinks that individual goodness measures are ill-defined. So even if we agree that the risk-free interpretation correctly expresses the core of egalitarianism, we are left in the dark about where this principle takes issue with the premisses of the best defence of utilitarianism, and why.

We will return to the topic of individual goodness measures in section 19, which elaborates on the reasons for not taking them as given in thinking about the ethics of distribution. But for now I have only argued that the risk-free interpretation faces a serious worry about arbitrariness, and leaves basic questions about what is at issue between utilitarianism and egalitarianism unanswered. Despite the naturalness of the risk-free interpretation, I therefore prefer to develop the fundamental egalitarian idea without presupposing the existence of individual goodness measures. Later arguments will provide further criticisms of the risk-free interpretation.

4. Terminology

Instead of working with individual goodness measures, I will be talking about risk. This section reviews some terminology to do with risk, and introduces some ideas to do with expected utility theory, including the important concept of a preorder. My account of egalitarianism will not depend on the use of expected utility theory, but I will use it for concrete examples.

The *betterness relation* holds between two lotteries just in case the first is at least as good as the second. The general form of lotteries is $[p_1, h_1; \dots; p_m, h_m]$ with the h_j 's the histories which could result from the lottery and the p_j 's their corresponding probabilities, all positive and adding up to one. The *risk-free betterness relation* holds between two histories just in case the first is at least as good as the second. Thus treating histories as lotteries which give the history probability one, the risk-free betterness relation is a special case of the betterness relation.

Now it is a very good question where these probabilities are coming from. In particular, how should we think of uncertainty when doing

ethics, and why should we think it representable in terms of probabilities? For some basic thoughts on this, see McCarthy forthcoming a. But the brief answer is that this is a difficult and important problem which we are a long way from fully understanding. We could make partial progress by showing that various views in ethics could be formulated without presupposing that, for example, uncertainty is probabilistically representable. I am optimistic that this could be done for the account of egalitarianism to be developed here, but I am not going to attempt it. So when I say that I will be offering an account of egalitarianism which essentially involves risk, I am thinking of risk as a placeholder for uncertainty. The assumption that probabilities are somehow supplied is just a simplification which enables me to dodge a whole cluster of difficult problems.

An individual i 's *individual betterness relation* holds between two lotteries just in case the first lottery is at least as good for i as the second. In the illustrative parts of the discussion I will typically assume that individual betterness relations satisfy the expected utility axioms (for which see, for example, Resnik 1987, Kreps 1988). That implies that for any individual i , there is a real-valued function u_i defined on histories such that the real-valued function U_i , defined on lotteries, represents i 's individual betterness relation, where for any lottery $L = [p_1, h_1; \dots; p_m, h_m]$,

$$(1) \quad U_i(L) := p_1 u_i(h_1) + \dots + p_m u_i(h_m)$$

Thus $U_i(L)$ is the expected value of u_i under the lottery L . The function u_i is unique up to positive affine transformation. In other words, suppose that v_i is some other function defined on histories. Then V_i , defined analogously to U_i , also represents i 's individual betterness relation if and only if $v_i = au_i + b$ for some real numbers a and b , with $a > 0$. Here and elsewhere I am using uppercase and lowercase letters to denote things to do with lotteries and histories respectively.

In the illustrative parts of the discussion I will also typically assume that interpersonal comparisons are unproblematic, in the sense that for any individuals i and j and any lotteries L_1 and L_2 , either L_1 is at least as good for i as L_2 is for j , or L_2 is at least as good for j as L_1 is for i . This means that we can pick functions u_i and U_i satisfying (1) for each individual i such that for any individuals i and j (not necessarily distinct) and any lotteries L_1 and L_2 , L_1 is at least as good for i as L_2 is for j iff $U_i(L_1) \geq U_j(L_2)$.

We might as well go one step further. Say that the *extended individual betterness relation* holds between individual/lottery pairs (i, L_1) and (j, L_2) just in case L_1 is at least as good for i as L_2 is for j . This essentially just rolls both interpersonal and intrapersonal comparisons into a single relation. Define a function U on individual/lottery pairs by $U(i, L) := U_i(L)$. Then U represents the extended individual betterness relation. Let u be the restriction of U to individual/history pairs. Then u represents the risk-free extended individual betterness relation.

I will call any function U which satisfies all these properties an *expectational function* which represents the extended individual betterness relation, and will summarize the assumptions which guarantee its existence via the claim that the extended individual betterness relation satisfies the expected utility axioms. I will sometimes call the function u which results from restricting any such U to individual/history pairs an *individual utility function*. This kind of construction is well known (for details see, for example, Hammond 1991, Broome 2004). From now on, U will always be some fixed expectational function which represents the extended individual betterness relation, and it will provide the numbers which appear in examples. Thus the *utility profile* of a history h will be understood as the n -tuple of real numbers $[u_1(h), \dots, u_n(h)]$. The utility profile of a lottery is defined by replacing each history which could result from it with its utility profile.

The basics of expected utility theory are by now common currency, and I will be using it in two ways. First, I will often use numbers in examples, in particular by describing histories as utility profiles. But in the use I make of the examples, the only significance of the numbers is that two lives with the same number are equally good, and a life with a greater number is better than a life with a lesser number. Here the use of expected utility theory is completely inessential, and I only use it because of its familiarity. Second, in discussing various specific theories, such as specific forms of egalitarianism, I will generally work within the expected utility framework. This is a useful simplifying assumption, enabling me to present straightforward results and focus on the philosophy, rather than delve into the mathematics needed to handle a more general framework.

Nevertheless, there are two crucial respects in which I will not be using expected utility theory. First, having renounced the use of individual goodness measures, I am not about to quietly start using individual utility functions as surrogates. For example, I will not be presupposing that a history with a utility profile $[1, 1]$ is, according to egalitarianism, better than one with profile $[2, 0]$ just because it

contains the same sum of individual utility but is more equal. Second, my central claims about the core of egalitarianism will make no use at all of expected utility theory. They will make sense in a framework in which just about every axiom of expected utility theory is abandoned. By contrast, some discussions of egalitarianism make heavy use of expected utility theory, leaving it unclear how the claims made about egalitarianism could be decoupled from the expected utility framework. This is important, because the claim that various evaluative relations satisfy the expected utility axioms is controversial in ways which seem to have little to do with egalitarianism. Thus not only do such discussions foist assumptions which seem to be needlessly strong upon egalitarianism, but they actually seem to obscure its core.

Now this is not the place to rehearse the ways in which the expected utility axioms are controversial. But one of the controversies can be illustrated with a concept which will play an important role later on. A *preorder* is a binary relation which is reflexive and transitive. To recall, if R is a binary relation on a set X , it is *reflexive* if for all x in X , xRx . Thus an ordering is a complete preorder. When being formal, variants on ' \succeq ' will always denote preorders. As is usual, $x \succ y$ is then defined as $x \succeq y$ but not $y \succeq x$, while $x \sim y$ is defined as $x \succeq y$ and $y \succeq x$. For example, if ' \succeq ' is interpreted as 'at least as good as', ' \succ ' means 'better than' and ' \sim ' means 'exactly as good as'.

Now one of the expected utility axioms is that the binary relation in question is an ordering. But in the case where the binary relation is some sort of evaluative comparative, this may not seem very plausible. For example, the assumption that the extended individual betterness relation is complete implies that for any histories h_1 and h_2 and individuals i and j , either h_1 is at least as good for i as h_2 is for j , or vice versa. Or more informally, for any two lives, either one is better than the other, or they are equally good.

However, many, such as Rawls (1982), will see this as implausible. I therefore wish to avoid assuming that evaluative comparatives are orderings. But it is one thing to wish to avoid that assumption, another to know how exactly to do that. The most obvious idea is to abandon completeness, and to assume only that the evaluative comparatives in question are preorders. This means they are allowed (but not required) to have gaps, making room for the claim that there can be two lives which are so different that it is neither the case that one is a better life than the other, nor the case that they are equally good.

I do not claim that retreating to preorders is the best way of dealing with worries about incomparability. But the general topic is very complicated, and preorders at least provide a relatively simple model. Thus my account of the core of egalitarianism will only assume that the extended individual betterness relation is a preorder and will not assume that it satisfies any of the other expected utility axioms.

5. The risk-involving interpretation

This section gives a schematic statement of what I take the core of egalitarianism to be. Assume a population of two people, Alice and Bob, and that a fair coin is going to be tossed. Consider an example due originally, I believe, to Myerson (1981), and discussed further by Broome (1989, 1991, forthcoming), Ben-Porath et al. (1997), and Fleurbaey (2010), among others.

L_E	<i>heads</i>	<i>tails</i>	L_F	<i>heads</i>	<i>tails</i>
<i>Alice</i>	1	0	<i>Alice</i>	1	0
<i>Bob</i>	1	0	<i>Bob</i>	0	1

The two lotteries L_E and L_F are equally good for each person. So from the point of view of a sole concern with well-being, they are equally good. But L_E guarantees equality while L_F guarantees inequality. So from the point of view of a sole concern with equality, L_E is uncontroversially better than L_F . Therefore, all egalitarians should think that L_E is on-balance better than L_F . But for later on, notice that L_F at least gives everyone a fair chance, hence the notation.

This suggests an alternative way of interpreting the fundamental egalitarian idea, based on the idea of dominance. Very roughly, suppose that \succeq_v is a preorder which compares parts of some class of objects with respect to some value V . Then X *weakly dominates* Y (in terms of the value V) if (i) X and Y can be expressed as the sum of nonoverlapping parts x_1, \dots, x_m and y_1, \dots, y_m respectfully such that for each j , x_j is in some sense the same size or importance as y_j ; and (ii) $x_j \succeq_v y_j$ for each j . X *strictly dominates* Y if X weakly dominates Y but Y does not weakly dominate X . X is *equivalent* to Y if X weakly dominates Y and Y weakly dominates X . The claim that egalitarians should judge that L_E is better than L_F then rests on two dominance ideas.

First, L_E and L_F are each the sum of two sublotteries, one involving heads and the other involving tails. The two involving heads have the same importance as they have the same probabilities, and likewise the

two involving tails. But $[1,1]$ is uncontroversially better in terms of equality than $[1,0]$, and $[0,0]$ is uncontroversially better in terms of equality than $[0,1]$. Therefore, L_E strictly dominates L_F in terms of minimal equality.

Second, L_E and L_F are each also the sum of two individual lotteries, the one Alice faces and the one Bob faces. These individual lotteries have the same importance in that they each involve a single individual. But the two individual lotteries involving Alice are equally good in terms of individual well-being, as are the two involving Bob. Therefore, L_E and L_F are equivalent in terms of well-being.

In brief, egalitarians care about exactly two things, well-being and equality. L_E is equivalent to L_F in terms of well-being, while it strictly dominates L_F in terms of minimal equality. Therefore, egalitarians should judge that L_E is better than L_F . In fact, this last inference could be seen as a third application of dominance. Thus Myerson's example suggests clause (a) of the following, while the other clauses have an analogous motivation.

The risk-involving interpretation of the fundamental egalitarian idea:

For any lotteries L_1 and L_2 , the following hold:

- (a) If L_1 is equivalent to L_2 in terms of well-being, and strictly dominates L_2 in terms of minimal equality, then L_1 is better than L_2
- (b) If L_1 is equivalent to L_2 in terms of well-being, and equivalent to L_2 in terms of minimal equality, then L_1 and L_2 are equally good
- (c) If L_1 strictly dominates L_2 in terms of well-being, and weakly dominates L_2 in terms of minimal equality, then L_1 is better than L_2

This idea has obvious analogies with the risk-free interpretation of the fundamental egalitarian idea. The disanalogies are that the risk-involving interpretation talks about risk, and in only using evaluative comparatives, makes no use of individual goodness measures.

The risk-involving interpretation is still rough. To make it more precise we will need accounts of weak dominance in terms of well-being and also in terms of minimal equality. To fully characterize egalitarianism, we will also need to supplement the risk-involving interpretation with various background ideas. We now turn to these topics.

6. Dominance in terms of well-being

One obvious proposal for interpreting dominance in terms of well-being is to say that a lottery L_1 weakly dominates L_2 in terms of well-being if for each member of the population i , L_1 is at least as good for i as L_2 . This is a Pareto-style condition.

However, consider two histories with utility profiles $[1, 0]$ and $[0, 1]$. From the point of view of a sole concern with well-being, the identities of individuals do not matter: Alice getting 1 and Bob getting 0 is as good as Alice getting 0 and Bob getting 1. So it is natural to judge the two histories to be equivalent in terms of well-being. The same point applies to lotteries. Thus dominance in terms of well-being is also naturally interpreted as satisfying an anonymity-style condition: if the only relevant difference between two lotteries lies in the identities of the individuals, the two lotteries are equivalent in terms of well-being.

Combining these two ideas suggests the following. A *permutation* of the population is a mapping σ of the population onto itself, so that every individual is paired off with a unique individual. Consider the following definition:

Definition 1

Let \succsim_p be the preorder on lotteries defined by: $L_1 \succsim_p L_2$ iff there is a permutation σ of the population such that for each member of the population i , L_1 is at least as good for i as L_2 is for $\sigma(i)$

We can think of \succsim_p as the *anonymous Pareto preorder*. The following is then natural.

Proposal 1

For any lotteries L_1 and L_2 , L_1 weakly dominates L_2 in terms of well-being iff $L_1 \succsim_p L_2$

This proposal makes no use of individual goodness measures, and is well-defined on the very weak assumption that the extended individual betterness relation is a preorder, with no further expected-utility-style assumptions.

7. Dominance in terms of minimal equality

My goal now will be to search for a concrete account of when one lottery weakly dominates another in terms of minimal equality. The challenging part of this task is to say what is the content of the minimal betterness in terms of equality relation. Earlier I said that this

relation holds between two histories h_1 and h_2 just in case h_1 is at least as good in terms of equality as h_2 according to any genuine form of egalitarianism. This would be fine as a definition of the minimal betterness in terms of equality relation if we had (a) an independent characterization of egalitarian accounts of the betterness relation, and (b) a method of extracting what any such account implicitly says about betterness in terms of equality. But since our project is to use the minimal betterness in terms of equality relation to characterize egalitarianism, this way of trying to secure the definition is a non-starter. This should leave us uneasy about whether there is indeed any such relation as the minimal betterness in terms of equality relation.

Nevertheless, it seems that there ought to be such a relation, and that characterizing it should be simpler than characterizing egalitarianism. For egalitarian accounts of overall betterness are surely some kind of attempt to combine concerns with well-being and concerns with equality, and it should be easier to characterize the basic components of the view than the overall picture. So in this section I will be offering a self-standing account of the content of the minimal betterness in terms of equality relation, and hoping that others will find it plausible.

In more detail, the minimal betterness in terms of equality relation will be a preorder between histories. Because different forms of egalitarianism disagree about when one history is better in terms of equality than another, the minimal betterness in terms of equality relation is bound to be incomplete. I will now discuss four constraints we might wish to impose upon it. I will argue that we ought to accept three of these, and then use the objection to the fourth to argue that there are unlikely to be any further constraints we should accept. Of course, having rejected the use of individual goodness measures, I will not be making any use of them, nor will I be using individual utility functions as surrogates.

7.1 Anonymity

Say that two histories h_1 and h_2 are *anonymously equivalent* just in case there is some permutation σ of individuals such that for each individual i , h_1 is exactly as good for i as h_2 is for $\sigma(i)$. Then a preorder of histories \succeq is *anonymous* just in case: for any anonymously equivalent histories h_1 and h_2 , $h_1 \sim h_2$. So roughly speaking, a preorder of histories is anonymous just in case it is indifferent to the identities of individuals. We should obviously require the minimal betterness in terms of equality relation to be anonymous. The real question is what requirements to do with equality we should impose upon it.

7.2 Equality

Say that a preorder of histories \succeq *favours equality* if for any histories h_1 and h_2 : (i) if h_1 and h_2 each contain equality, then $h_1 \sim h_2$; and (ii) if h_1 contains equality and h_2 does not, then $h_1 \succ h_2$. We have to require the minimal betterness in terms of equality relation to favour equality. But it is a weak requirement. Anyone concerned with equality cares not just about the elimination of inequality, but also its reduction. So the next question is what the minimal betterness in terms of equality relation should say about the reduction of inequality.

7.3 Local reductions

Some terminology will help. Say that two individuals are *strictly closer together* in h_1 than in h_2 if: (i) in h_2 , one of them is better off than the other — call the former i and the latter j ; (ii) i is worse off in h_1 than in h_2 ; (iii) j is better off in h_1 than in h_2 ; and (iv) in h_1 , i is still at least as well off as j .

Say that h_1 is *strictly locally more equal* than h_2 if there are two individuals j and k such that (i) for every individual i apart from j and k , h_1 is exactly as good for i as h_2 ; and (ii) j and k are strictly closer together in h_1 than in h_2 . For example, $[0, 2, 3, 5]$ is strictly locally more equal than $[0, 1, 4, 5]$ because individuals 2 and 3 are strictly closer together in the first while individuals 1 and 4 are unaffected. Say that a preorder \succeq *favours strict local reductions in inequality* just in case for any histories h_1 and h_2 , if h_1 is strictly locally more equal than h_2 , then $h_1 \succ h_2$.

There are three reasons for taking seriously the idea of requiring the minimal betterness in terms of equality relation to favour strict local reductions in inequality. First, ignoring a minor difference, the idea that strict local reductions in inequality are always improvements is what the so-called Hammond equity principle says (Hammond 1976). Later writers (e.g. Tungodden 2003) have taken the Hammond equity principle to be an extreme egalitarian idea. Thus there is tacit support in the literature for the idea that strict local reductions in inequality are at least always improvements in terms of equality (whether or not they are always on balance improvements).

Second, a so-called Pigou-Dalton transfer of individual goodness takes a unit of individual goodness from a better-off person i and gives it to a worse-off person j while still leaving i at least as well off as j . The *Pigou-Dalton principle* then says that if a history h_1 can be obtained from a history h_2 by a Pigou-Dalton transfer of individual goodness, then h_1 is better than h_2 . Now the Pigou-Dalton principle is often said

to be a basic egalitarian idea, giving concrete expression to the idea that equality is valuable. But as a basic egalitarian idea, it is expressed within a framework this article has rejected, one which takes the existence of individual goodness measures for granted. Nevertheless, we can still extract something useful from it. The Pigou-Dalton principle is about overall betterness. But a natural way of decomposing it is into part (a) of the risk-free interpretation of the fundamental egalitarian idea together with the view that strict local reductions are always improvements in terms of equality. But of course, the latter idea can be detached from the framework of the risk-free interpretation, and can be accepted without taking the existence of individual goodness measures for granted.

Third, requiring the minimal betterness in terms of equality relation to favour strict local reductions in inequality is tempting in part because because one can turn a history with inequality into one with perfect equality by finitely many strict local reductions in inequality, so it may seem to be a natural extension of the idea that the minimal betterness in terms of equality relation favours equality.

Nevertheless, we should not require the minimal betterness in terms of equality relation to favour strict local reductions in inequality. Consider this example. Let $h_1 = [10, 10, 10, 10, 10, 0, 0, 0, 0, 0]$ and $h_2 = [10, 10, 10, 10, 6, 4, 0, 0, 0, 0]$. Obviously, h_2 is strictly more locally equal than h_1 . Some egalitarians may think that h_2 is better in terms of equality than h_1 . But other egalitarians may think that h_2 is worse in terms of equality than h_1 . For h_2 contains more social divisions; each equally well-off subgroup has fewer members; and some individuals are entirely isolated. Such an egalitarian might think that h_1 is doing pretty well in terms of equality, for each individual is a member of a fairly large group within which there is perfect equality. Such an egalitarian might also think that to the extent that someone thinks that h_2 is better in some important way than h_1 , she is better understood as expressing the kind of intuition Parfit (2000) claimed underlies the priority view, that there is some kind of nonegalitarian rationale for giving greater weight to improving the position of the worse off.

Now it is not my purpose to assess the relative plausibility of these two opposing views about h_1 and h_2 . I only wish to claim that the view which says that h_2 is worse in terms of equality than h_1 articulates recognizably egalitarian ideas in an intelligible way, and should clearly count as an egalitarian position. We should therefore not require the minimal betterness in terms of equality relation to favour strict local reductions in inequality.

On reflection, this is not surprising. Moving two people strictly closer would be an improvement in terms of equality if they were the only people around. But when others are around, moving the two strictly closer while leaving everyone else unchanged is sometimes going to result in one of them being moved further away from others. In such cases, one aspect of the change will be a local improvement in inequality, and another will be a local worsening in inequality. The proposal that the improvement will uncontroversially always outweigh the worsening is far too strong. So we now seek to weaken the proposal while preserving what is appealing about it, and an idea due to Vallentyne (2000) seems to provide the solution.

7.4 Contractions

Roughly speaking, a contraction brings the endpoints of two distributions (the worst-off and the best-off) closer together, without changing anything else. Thus the worst-off are still among the worst-off, and the best-off are still among the best-off; they are just not so far apart. For example, $[1, 1, 3, 3]$ is a contraction of $[0, 1, 4, 4]$. Notice that the best-off group went down together in moving from $[0, 1, 4, 4]$ to $[1, 1, 3, 3]$, while the worst-off group gained a new member.

It is easy to formalize the idea of a contraction when the risk-free extended individual betterness relation is complete. Ignoring a small slip, this is done in Vallentyne 2000; Tungodden 2000, 2003; Tungodden and Vallentyne 2005. But because I wish my account of egalitarianism to cater for incompleteness, the task is a bit more tricky. The basic problem is that in the absence of completeness, there may not be a best-off (or worst-off) group. To illustrate, suppose (implausibly) that there are two totally incomparable goods, and that one person is better off than another iff she has at least as much of one good and more of the other. We can denote each person's holding of these goods by a pair of numbers $\langle x, y \rangle$ where the only significance of these numbers is that someone with $\langle x_1, y_1 \rangle$ is at least as well off as someone with $\langle x_2, y_2 \rangle$ if and only if $x_1 \geq x_2$ and $y_1 \geq y_2$. Because the goods are totally incomparable, comparisons between the x 's and y 's have no meaning. Thus it is not implied that an individual has the same level of each good in $\langle 1, 1 \rangle$, or that in moving from $\langle 0, 1 \rangle$ to $\langle 1, 2 \rangle$ her holding of each good increases by the same amount.

Consider then a history $h_1 = [\langle 0, 0 \rangle, \langle 1, 2 \rangle, \langle 2, 1 \rangle]$ containing three people, A, B, and C. Then there is no individual who is best off, in the sense of being at least as well off as every other individual. Nevertheless, $h_2 = [\langle 0, 0 \rangle, \langle 1, 1 \rangle, \langle 1, 1 \rangle]$ is intuitively a contraction of h_1 , and an

unequivocal improvement in equality. Notice that in h_1 there are two maximally well-off individuals, B and C. In moving from h_1 to h_2 , the maximally well-off individuals have moved down to become part of the best-off group, which is now nonempty. By contrast, moving from h_1 to $h_3 = [(0, 0), (1, 1), (2, 1)]$ is intuitively not a contraction. The definitions which follow are intended to generalize these ideas.

For any history h , let $Best(h)$ be the set of individuals in h , possibly empty, who are at least as well off as everyone. Similarly, let $Worst(h)$ be the set of individuals such that everyone is at least as well off as them. So roughly speaking, $Best(h)$ and $Worst(h)$ denote the best-off and worst-off individuals in h respectively, except that there may not be any such individuals. All individuals in $Best(h)$ are equally well off, as are all individuals in $Worst(h)$. Let $Max(h)$ be the set of individuals in h who are no worse off than anyone, and let $Min(h)$ be the set of individuals in h who are no better off than anyone. These sets can never be empty, and $Best(h) \subseteq Max(h)$ along with $Worst(h) \subseteq Min(h)$. For example, $Best(h_1) = \emptyset$, the empty set, $Max(h_1) = \{B, C\}$, $Worst(h_1) = \{A\}$ and $Min(h_1) = \{A\}$.

Definition 2

A history h_1 is a *downwards contraction* of a history h_2 if and only if

- (i) h_2 contains inequality
- (ii) for every individual i in $Max(h_2)$, i is worse off in h_1 than in h_2
- (iii) $Max(h_2) \subseteq Best(h_1)$
- (iv) for every individual i not in $Max(h_2)$, i is exactly as well off in h_1 as in h_2

In the special case in which the risk-free extended individual betterness relation is complete (i.e. there are no problems with inter- or intrapersonal comparisons), this just says that a downwards contraction moves the best-off down, but not below the next best-off group, and leaves everyone else unchanged. We likewise have

Definition 3

A history h_1 is an *upwards contraction* of a history h_2 if and only if

- (i) h_2 contains inequality
- (ii) for every individual i in $Min(h_2)$, i is better off in h_1 than in h_2
- (iii) $Min(h_2) \subseteq Worst(h_1)$

- (iv) for every individual i not in $Min(h_2)$, i is exactly as well off in h_1 as in h_2

Likewise, when there are no problems with inter- or intrapersonal comparisons, an upwards contraction moves the worst-off up, but not above the next worst-off group, and leaves everyone else unchanged.

Now, what to say about contractions in the face of incompleteness deserves more attention than I have given it. But a serious discussion of incompleteness should really be the topic of a separate article. Here I am trying to say just enough, and none of my later arguments will rely on the particular way I have defined contractions in the face of incompleteness. So setting this issue aside, the fundamental concept we have been seeking is as follows:

Definition 4

A history h_1 is a *contraction* of a history h_2 just in case h_1 can be obtained from h_2 by either a downwards contraction, an upwards contraction, or a downwards and an upwards contraction

Say that a preorder \succeq *favours contractions* just in case for any histories h_1 and h_2 , if h_1 is a contraction of h_2 , then $h_1 \succ h_2$. Contractions unequivocally reduce inequality without always eliminating it, but they never give rise to the worry which led to the rejection of the claim that the minimal betterness in terms of equality relation favours strict local reductions in inequality. We should accept the weaker claim that the minimal betterness in terms of equality relation favours contractions.

7.5 Conclusion

Should we impose further conditions on the minimal betterness in terms of equality relation? I believe not. It would be tedious to enumerate the possibilities precisely, but two examples should make the grounds for scepticism clear. For simplicity, we assume the risk-free extended individual betterness relation is complete.

In history h_1 one person is well off and the other is very well off. In h_2 one is badly off and the other is very badly off. There is inequality in both histories, but which is worse in terms of equality? On one view, equality matters more when people are well off (a luxury for the rich). On another view, equality matters more when people are badly off (a panacea for the poor). I am not concerned with the relative plausibility of these views. As Sen (1973) points out, both views are found in the literature. We should therefore conclude that h_1 and h_2 are

incomparable with respect to the minimal betterness in terms of equality relation.

Say that a history h_1 is *bounded by* h_2 just in case the best-off in h_2 are at least as well off as the best-off in h_1 , and the worst-off in h_1 are at least as well off as the worst-off in h_2 . Suppose that neither of h_1 and h_2 is bounded by the other, and that there is equality in neither. Then because there are different attitudes to how the badness of inequality varies with how well-off people are, I believe we should conclude that h_1 and h_2 are incomparable with respect to the minimal betterness in terms of equality relation. In the case where one of h_1 and h_2 is bounded by the other, but neither can be obtained from the other by a series of contractions, I believe we should again conclude that h_1 and h_2 are incomparable with respect to the minimal betterness in terms of equality relation. For without going into details, in this case the transition from h_1 to h_2 is going to involve local changes in inequality, some for the better and others for the worse, and the objection to requiring the minimal betterness in terms of equality relation to favour strict local reductions in inequality is going to recur.

In summary, I believe we should require the minimal betterness in terms of equality relation to be anonymous and to favour equality and contractions, but we should not require it to satisfy any further conditions about equality. Towards defining it, and preparing the ground for some other definitions, consider the following definition:

Definition 5

- (i) \succsim_w is the smallest anonymous preorder of histories which favours equality; (ii) \succsim_e is the smallest anonymous preorder of histories which favours both equality and contractions; and (iii) \succsim_l is the smallest anonymous preorder of histories which favours equality, contractions, and strict local reductions

These preorders are well-defined.² What follows is the main claim of this section.

Proposal 2

The minimal betterness in terms of equality relation is \succsim_e

The argument that there is such a relation as the minimal betterness in terms of equality relation just rests on the naturalness of Proposal 2.

² Just to illustrate one case, define a binary relation E between histories such that Eh_1h_2 iff h_1 and h_2 are anonymously equivalent, or h_1 contains equality, or h_1 is a contraction of h_2 . Then \succsim_e is the so-called transitive closure of E .

Corresponding to the preorders in Definition 5 are three dominance relations.

Definition 6

For any lotteries of the form $L_1 = [p_1, h_1; \dots; p_m, h_m]$ and $L_2 = [p_1, k_1; \dots; p_m, k_m]$, say that: (i) $L_1 \succsim_W L_2$ iff $h_j \succsim_w k_j$ for every $j = 1, \dots, m$; (ii) $L_1 \succsim_E L_2$ iff $h_j \succsim_e k_j$ for every j ; and (iii) $L_1 \succsim_L L_2$ iff $h_j \succsim_l k_j$ for every j

Since \succsim_w , \succsim_e and \succsim_l are preorders of histories, \succsim_W , \succsim_E and \succsim_L are preorders of lotteries. Then Proposal 2 yields

Proposal 3

For any lotteries L_1 and L_2 , L_1 weakly dominates L_2 in terms of minimal equality just in case $L_1 \succsim_E L_2$

8. Characterization of egalitarianism

This section begins with an account of the core of egalitarianism. It then discusses how to embed this core into a full-blown characterization of egalitarianism. Throughout this section, \succsim is the betterness relation.

8.1 The core

The previous two sections offered accounts of ‘weakly dominates in terms of well-being’ and ‘weakly dominates in terms of minimal equality’. We have just defined the latter via Proposal 3. The former was defined in section 6 via Proposal 1, which said that a lottery L_1 weakly dominates L_2 in terms of well-being just in case $L_1 \succsim_P L_2$, where \succsim_P is the anonymous Pareto preorder. Thus slotting Proposals 1 and 3 into the risk-involving interpretation of the fundamental egalitarian idea yields what I suggest is the core of egalitarianism:

The core of egalitarianism:

For any lotteries L_1 and L_2 the following hold:

- (a) $L_1 \sim_P L_2$ and $L_1 \succ_E L_2 \Rightarrow L_1 \succ L_2$
- (b) $L_1 \sim_P L_2$ and $L_1 \sim_E L_2 \Rightarrow L_1 \sim L_2$
- (c) $L_1 \succ_P L_2$ and $L_1 \succsim_E L_2 \Rightarrow L_1 \succ L_2$

My overall account of the core of egalitarianism therefore rests on four components: (i) the fundamental egalitarian idea; (ii) the risk-involving interpretation as a way of making the fundamental egalitarian

idea more concrete; (iii) a specific account of when one lottery dominates another in terms of well-being; and (iv) a specific account of when one lottery dominates another in terms of minimal equality.

This account makes good on the goal set in section 3 of developing an account of the core of egalitarianism without appealing to individual goodness measures. The account also has two attractive theoretical properties, stability and generality.

I will discuss stability in section 9.2, but generality comes from the fact that the account makes few assumptions about the extended individual betterness relation and the betterness relation. It merely assumes that those relations are preorders. It does not assume, for example, that individual betterness relations satisfy any of the further axioms of expected utility theory, namely completeness, strong independence, and continuity. As already noted, many discussions of egalitarianism do make such assumptions, and it is unclear how their accounts of egalitarianism could be decoupled from them. Yet the assumptions are controversial, so it is better for an account of egalitarianism not to have to rely on them.

8.2 *Theory E*

We have now filled in the internal details of the risk-involving interpretation. But we earlier noted that to fully characterize egalitarianism, the risk-involving interpretation would need to be supplemented with various background ideas. I will now discuss a natural way of doing that.

Recall that utilitarianism is said to be indifferent to equality, and that I am taking the best defence of utilitarianism to be based on a variant of Harsanyi's aggregation theorem. It is not obvious what, if anything, in the premisses of Harsanyi's theorem conflicts with the idea that equality is valuable. But suppose there is a conflict, and we manage to isolate it. Then the other premisses will be common ground between utilitarianism and egalitarianism, and should provide the needed background ideas.

To pursue this proposal, I need to state the premisses of the Harsanyi-like result which provides the best defence of utilitarianism. The premisses of Harsanyi's original theorem undoubtedly contain bells and whistles: ideas which are not essential to utilitarianism, but which help make the theorem easier to present and prove; in brief, useful simplifications and idealizations. For example, the original theorem assumes that the various evaluative relations in play are

complete; but as already mentioned at the end of section 4, that is a bit hard to believe.

So to properly pursue the proposal, I would need to state a Harsanyi-like theorem which is free from bells and whistles, and defend the claim that this theorem captured everything which was essential to utilitarianism. I believe this can be done, but I do not have the space to discuss it.³ Instead, I am just going to rehearse a version of Harsanyi's original theorem, bells and whistles included. Thus the eventual characterization of egalitarianism will itself include bells and whistles, and the task of paring those down will be left for another day.

There are several ways of presenting the Harsanyi premisses, but consider the following:

The anonymous Pareto principle:

For any lotteries L_1 and L_2 , the following hold:

(a) $L_1 \sim_P L_2 \Rightarrow L_1 \sim L_2$

(b) $L_1 \succ_P L_2 \Rightarrow L_1 \succ L_2$

Minus the 'anonymous' bit, this is what Broome (1991) calls the principle of personal good. Then these are one way of putting the premisses of Harsanyi's theorem.

The utilitarian premisses:

- (a) The extended individual betterness relation satisfies the expected utility axioms
- (b) The betterness relation satisfies the expected utility axioms
- (c) The anonymous Pareto principle

With mild domain assumptions, Harsanyi was able to show that if the utilitarian premisses are true, the betterness relation is represented by the function S defined on lotteries, where for any lottery L , $S(L) := U_1(L) + \dots + U_n(L)$. So one lottery is better than another iff it contains a greater sum of individual expected utilities, or equivalently, a greater expected sum of individual utilities.

The only utilitarian premiss which conflicts with (more precisely, is inconsistent with) the egalitarian core is anonymous Pareto. But this is a crude account of the conflict, and we can be more refined. The egalitarian core is the conjunction of the following two ideas.

³ What I think is the needed theorem is presented in McCarthy et al. MS and its interpretation is discussed in McCarthy forthcoming b.

The equality principle:

For any lotteries L_1 and L_2 : $L_1 \sim_P L_2$ and $L_1 \succ_E L_2 \Rightarrow L_1 \succ L_2$

The equality-neutral anonymous Pareto principle:

For any lotteries L_1 and L_2 :

$$(a) L_1 \sim_P L_2 \text{ and } L_1 \sim_E L_2 \Rightarrow L_1 \sim L_2$$

$$(b) L_1 \succ_P L_2 \text{ and } L_1 \succsim_E L_2 \Rightarrow L_1 \succ L_2$$

Three points are obvious, but worth making explicit. First, the equality principle conflicts with anonymous Pareto. For example, in Myerson's example, the equality principle implies that L_E is better than L_F , while anonymous Pareto implies that they are equally good. Second, anonymous Pareto is logically strictly stronger than equality-neutral anonymous Pareto. Roughly, (a) of anonymous Pareto says that nothing matters aside from well-being, and (b) says that well-being does matter. Equality-neutral anonymous Pareto agrees with this in the special case where considerations of equality cannot conflict, but is otherwise silent. Third, there is no conflict between equality-neutral anonymous Pareto and the equality principle.

Therefore, the natural way of weakening the utilitarian premisses to find the common ground between utilitarianism and egalitarianism is to start with the utilitarian premisses then weaken them by replacing anonymous Pareto with the weaker equality-neutral anonymous Pareto. To obtain egalitarianism, we then add the equality principle. Thus the bells-and-whistles characterization of egalitarianism is as follows:

The egalitarian premisses:

- (a) The extended individual betterness relation satisfies the expected utility axioms
- (b) The betterness relation satisfies the expected utility axioms
- (c) The equality principle
- (d) The equality-neutral anonymous Pareto principle

The bells and whistles belong to the first two premisses, so one can get a rough sense of what a characterization of egalitarianism free from such idealizations would look like. I will now mostly suppress the qualification about bells and whistles.

There is much to do before I can complete the argument for the claim that the premisses listed above in fact are the premisses

(or axioms) of egalitarianism. In order not to continually beg the question, I need another name for the theory which has those premisses. I will call it theory E. It is consistent. Consider the following claim:

(E) Egalitarianism is theory E

Ignoring the point about bells and whistles, (E) is the main claim of this article.

9. Variations

Most egalitarians do not just claim that egalitarianism is correct; they also argue for particular forms of egalitarianism. So before turning to the more detailed defence of (E), I will briefly discuss some of the ways in which theory E can be strengthened. Since it can be strengthened in different ways, this means looking at what different forms of egalitarianism can disagree about. This discussion will also illustrate how my basic account of egalitarianism is stable under various revisions of detail.

9.1 *Betterness in terms of equality*

Theory E is built on top of what I claimed was the minimal betterness in terms of equality relation, \succsim_e . This was supposed to characterize when one history is better than another in terms of equality according to all genuine forms of egalitarianism. This leaves room for particular forms of egalitarianism to say more about betterness in terms of equality.

One natural type of egalitarian theory arises from extending \succsim_e to some richer preorder \succsim_q and claiming that although \succsim_e is the minimal betterness in terms of equality relation, \succsim_q captures the full account of when one history is better in terms of equality than another.⁴ Then we can build up particular egalitarian theories using \succsim_q as the account of betterness in terms of equality in the same way that we developed our account of egalitarianism itself from \succsim_e .

To illustrate, section 7 introduced the idea that strict local reductions of inequality are always improvements in terms of equality. I claimed that this idea is too strong to be built into a characterization of egalitarianism. But I also argued that at least implicitly, it is a fairly

⁴ \succsim_q is an extension of \succsim_e just in case for all histories h_1 and h_2 : (i) $h_1 \succsim_e h_2 \Rightarrow h_1 \succsim_q h_2$; and (ii) $h_1 \succ_e h_2 \Rightarrow h_1 \succ_q h_2$. In set-theoretic terms, $\succsim_e \subseteq \succsim_q$ and $\succ_e \subseteq \succ_q$.

popular egalitarian idea. The betterness in terms of equality preorder of histories generated by this idea is the preorder \succeq_I , and the corresponding dominance preorder of lotteries is \succeq_L , as per Definitions 5 and 6 in section 7.5. Then the core of the egalitarian theory built on top of this idea is obtained from replacing occurrences of ' \succeq_E ' with instances of ' \succeq_L ' in the statement of the egalitarian core, and a bells-and-whistles version, which we might call theory L, results from doing the same to the axioms of theory E. Thus theory L is a natural way of trying to embed the idea that strict local reductions are always improvements in terms of equality into a concrete egalitarian framework.

9.2 *Stability*

Some may disagree with aspects of my development of (E). But my basic account is stable under various revisions of detail.

First, someone might claim that the minimal betterness in terms of equality relation is not \succeq_e , as I claimed, but rather some different preorder \succeq_q . But just as I built theory E on top of \succeq_e , this critic could build theory Q on top of \succeq_q , then claim that egalitarianism is in fact theory Q. This critic disagrees with one of the details I offered, but the basic story remains the same.

To illustrate, some of the literature has tacitly claimed that \succeq_w (see Definition 5 in Sect. 7.5) is the minimal betterness in terms of equality relation, and could thus claim that the corresponding theory W is egalitarianism. For example, theory W is the main theory discussed by Fleurbaey (2010). Theory W is logically strictly weaker than theory E, and because it effectively subtracts the idea that contractions are always improvements from my statement of the egalitarian core, I think it is too weak to count as egalitarianism. But my main point is that someone could accept my general approach to egalitarianism while substituting their own account of some of the details, such as the content of the minimal betterness in terms of equality relation.

Second, another departure from the detail of my account would pack structural conditions in earlier on. The fundamental egalitarian idea talks about betterness in terms of well-being and betterness in terms of equality. The risk-involving interpretation interpreted those ideas in terms of the dominance preorders \succeq_P and \succeq_E . But one might well claim that betterness in terms of well-being and betterness in terms of equality should in addition be interpreted as satisfying some further expected utility axioms. Depending on which further axioms are added, this would result in a richer interpretation of both relations. By contrast, I waited until I discussed the betterness

relation to build in those kinds of background or structural conditions.

I am not going to pursue this because it would involve taking a much closer look at expected utility theory, and that would go beyond the scope of this article. Nevertheless, I am confident it makes no real difference. For example, once we accept an account of overall betterness which satisfies the expected utility axioms, as theory E does, adding those axioms to the account of betterness in terms of well-being and betterness in terms of equality turns out to be redundant: the resulting theory is equivalent to theory E. More generally, I believe it turns out to be largely a matter of taste whether one builds some fragment of the expected utility axioms directly into the accounts of betterness in terms of well-being and of equality, or whether instead one simply takes overall betterness to be governed by that fragment. I have chosen the latter because it simplifies the presentation.

9.3 Equality and well-being

Egalitarianism has two concerns, equality and well-being. It therefore seemed natural to start by developing accounts of these components separately, and then to arrive at an egalitarian account of overall betterness by examining how these two components interact. This follows the principle of trying to analyse wholes in terms of what would seem to be their simpler constituent parts.

However, many popular approaches to egalitarianism, particularly in economics, do not do this. Rather, they start with what they say are egalitarian claims about overall betterness, then derive results from those claims (see, for example, Ben-Porath et al. 1997, Harel et al. 2005). This raises the question of how such approaches can be understood in the present framework. In particular, it raises the question of whether such allegedly egalitarian claims about overall betterness are implicitly committed to a determinate view about betterness in terms of equality, and if so, whether there is a way of extracting that view from the claims about overall betterness.

Different forms of egalitarianism can disagree by disagreeing about betterness in terms of equality. But even if they agree about betterness in terms of equality, they can disagree about how to balance concerns with equality against concerns with well-being. It would help if we had some natural way of characterizing the tradeoffs which can be made between these concerns. A solution to this plus the problem of the previous paragraph is offered by McCarthy and Thomas (MS), but

here I am going to address two simpler topics about balancing well-being against equality.

9.4 Pareto

Speaking loosely, Pareto-style principles say that when there is some kind of overall improvement in well-being in which no one is made worse off, the result is an improvement in overall betterness.

When restricted to histories, theory E has fairly modest Pareto-style implications. For example, it follows from equality-neutral anonymous Pareto that if everyone is at least as well off in h_1 as in h_2 , and some are better off, and h_1 is at least as good in terms of equality as h_2 , then h_1 is better than h_2 . For example, it follows that $[1, 1]$ is better than $[1, 0]$. But it does not follow from theory E that, for example, $[1, 0]$ is better than $[0, 0]$.

This is as it should be. Theory E is intended to capture any egalitarian theory. Egalitarians care about both well-being and equality. Nothing follows from that about how to resolve a conflict between the two. Some resolutions will give more weight to well-being; others will give more weight to equality. One does not fail to be an egalitarian merely by giving more weight to equality, and thinking, for example, that $[1, 0]$ is on-balance worse than $[0, 0]$.

Nevertheless, most egalitarians think that $[1, 0]$ is better than $[0, 0]$. More generally, they accept that if one history is better for some than another, and at least as good for everyone, then it is better than the other. Parfit (2000) calls versions of egalitarianism which accept this condition *moderate*. More generally, most egalitarians will accept what I will call *risk-free anonymous Pareto*, which is just the restriction of anonymous Pareto to histories.

The assumption that there are moderate forms of egalitarianism has played an important role in the development of egalitarian theory. Parfit (2000) claimed that according to egalitarianism, it is in one way better if everyone is brought down to the level of the worst off, for that would achieve perfect equality. He suggested that this conclusion is absurd, and called it the *levelling-down* objection to egalitarianism. Some writers have responded by claiming that there are moderate versions of egalitarianism, and have then claimed that when applied to these versions of egalitarianism, the levelling-down objection is no real objection. Since these versions say that it would never be on balance better to make some worse off without making others better off, there is nothing absurd or repugnant about these

versions claiming that it would be in one way better (see, for example, Temkin 1993, p. 282, and Fleurbaey forthcoming).

I will return to this topic in the last two sections where I will argue that levelling down is a much more serious problem for egalitarians than has been acknowledged. There are other interesting Pareto-style conditions, but for now I have just wanted to illustrate another kind of way in which theory E can be augmented in a direction most egalitarians find appealing.

9.5 Extreme egalitarianism

Let me mention a stronger form of egalitarianism. This accepts that whenever there is any conflict between well-being and equality, equality always wins. Thus the core of extreme egalitarianism is obtained by adding to the core of egalitarianism the following claim: $L_1 \succ_E L_2 \Rightarrow L_1$ is better than L_2 . Thus extreme egalitarianism regards $[1, 1]$ as better than $[11, 10]$ (equality is better than inequality), and $[11, 10]$ as better than $[100, 10]$ (contractions are always improvements). But it does not disregard well-being, as it also regards $[1, 1]$ as better than $[0, 0]$. It is a strong theory, but that does not mean it is not a form of egalitarianism. It is similar to the account Rawls (1971) offered for the just distribution of some of the primary goods.

9.6 Variable populations

Finally, theory E does not address variable-population problems. But egalitarians will want to say something about these, and in the context of the present approach, two topics are worth noting. First, a fully general account of egalitarianism should say how expanding the population affects inequality. In particular, it should extend the minimal betterness in terms of equality relation to compare histories with different-sized populations. Second, at the very least, it should take into account the fact that Myerson's example has a variable-population analogue. Following Broome (2004), let us denote non-existence by Ω , so that in the following lotteries, Alice will exist in both states of nature, whereas Bob will only exist in one.

$L_{E'}$	<i>heads</i>	<i>tails</i>	$L_{F'}$	<i>heads</i>	<i>tails</i>
<i>Alice</i>	1	0	<i>Alice</i>	1	0
<i>Bob</i>	1	Ω	<i>Bob</i>	Ω	1

Egalitarians should clearly think that $L_{E'}$ is on-balance better than $L_{F'}$, so a fully general account of egalitarianism is going to have to capture

this kind of case. One possibility is that this will have no impact on (E): theory E is still the correct account of egalitarianism for constant populations. But another possibility is that variable-population considerations will impose additional constraints on the constant population case, leading to some stronger theory E* as the correct account of egalitarianism for constant populations. But variable-population problems need separate discussion, so this article will set this issue aside. However, even if theory E needs to be strengthened, the basic approach taken here, and almost all of the substantive claims, are unlikely to be affected.

10. Interim summary

My main claim is (E). So far I have argued that (E) emerges naturally from the risk-involving interpretation of the fundamental egalitarian idea. I have also argued that the basics of this account of egalitarianism are stable under revisions of detail, and have sketched how distinctions between different forms of egalitarianism can be expressed within my framework.

We will now apply (E) to a cluster of related topics. These include the way egalitarianism is concerned with patterns between people; how egalitarianism relates to other distributive views; what egalitarianism says about the distribution of individual goodness; and whether egalitarianism is true.

However, this discussion will as much be a further defence of (E) as an application of it. For example, there are popular views about some of the topics just mentioned which do not obviously harmonize with (E), and do not talk about risk. So these views form an objection to (E): (E) generates unsatisfactory positions on the topics in question (either wrong, or too indirect), and in talking about risk, (E) adds needless complexity. The reply will be that the popular views are mistaken, and part of the argument for (E) will be that it gives a better account, thanks to talking about risk.

We will start by looking closely at the way in which egalitarianism is concerned with patterns between people. Aside from the fact that many arguments in this area need to be unpacked, a detailed look at patterns will be important for the resolution of most of the topics just mentioned.

Most of the discussion about egalitarianism and patterns between people has revolved around a concept known as strong separability.

The next four sections will examine arguments for and against the idea that egalitarianism should partly be characterized in terms of the denial of strong separability. A diagnosis of problems with these arguments will then motivate a different take on the way egalitarianism is concerned with patterns, and what started out as an objection to (E) will end up redounding in its favour. The article will then apply the conclusion about patterns to other topics, including, in the last two sections, the truth of egalitarianism.

11. Incompatibility with strong separability?

Suppose two histories h_1 and h_2 are such that there is some subgroup of people such that for each member i of this subgroup, h_1 is exactly as good for i as h_2 . Thus members of the group are unaffected by the choice between the two histories. Then roughly speaking, the risk-free betterness relation is strongly separable across individual outcomes if and only if in any such case, we can ignore what h_1 and h_2 are like for members of the subgroup when we ask whether h_1 is at least as good as h_2 .

The view that egalitarianism has to deny strong separability has widespread support. Sen 1973, p. 41 can be read as expressing doubt about the compatibility of egalitarianism and strong separability: 'if one feels that the social valuation of the welfare of individuals should depend crucially on the levels of welfare ... of others' then strong separability 'has to be sacrificed'. Others go further. Broome says that '[i]f equality is valuable, strong separability has to go' (Broome 1989, p. 250), or the rejection of strong separability 'captures the egalitarian idea that comparisons between different people's wellbeing matter' (Broome forthcoming). And Temkin (1993, p. 138) says that strong separability obscures what is distinctive about egalitarian concerns: 'equality', he says, 'is an essentially comparative notion'.

Before giving an example to better explain what strong separability amounts to, here is a formal definition.

Definition 7

Suppose histories h_1 through h_4 are such that we can partition the population into two groups, I and J , such that (i) for all $i \in I$, h_1 is exactly as good for i as h_3 , and h_2 is exactly as good for i as h_4 , and (ii) for all $i \in J$, h_1 is exactly as good for i as h_2 , and h_3 is exactly as good for i as h_4 . Then a preorder \succeq on histories is *strongly separable across individual outcomes* iff in all such cases: $h_1 \succeq h_2$ iff $h_3 \succeq h_4$.

To save words, I will usually just refer to the claim that the risk-free betterness relation is strongly separable across individual outcomes as strong separability.

The only type of example I have seen to motivate the idea that egalitarians have to reject strong separability is an analogue of the Allais paradox (see Sen 1973 and Broome 1989), while the variant I will discuss is from Broome forthcoming. To present it, I will temporarily assume that we are working within the framework of the risk-free interpretation of the fundamental egalitarian idea, so that egalitarians are characterized as caring about two things: increasing the sum of individual goodness, and reducing the inequality of its distribution. Consider then the following histories, where the numbers are taken to measure individual goodness.

$$\begin{array}{ll} h_1 = [2, 2, 2, 2, 2, 2, 2, 2, 2, 2] & h_2 = [4, 1, 2, 2, 2, 2, 2, 2, 2, 2] \\ h_3 = [2, 2, 1, 1, 1, 1, 1, 1, 1, 1] & h_4 = [4, 1, 1, 1, 1, 1, 1, 1, 1, 1] \end{array}$$

Let I be composed of the first two members of the population, and J the rest. Then the definition of strong separability applies, with the conclusion that strong separability implies that h_1 is at least as good as h_2 iff h_3 is at least as good as h_4 . Informally, the members of J are unaffected by the choice between h_1 and h_2 , so strong separability implies that we can ignore them. The same is also true of h_3 and h_4 . But in each case, we are left with $[2, 2]$ versus $[4, 1]$. So if, for example, we think h_1 is better than h_2 , we should also think that h_3 is better than h_4 .

That just illustrates strong separability, so let us look at the case for saying that egalitarianism and strong separability are at least in tension. Making his terminology match mine, here is Sen (1973, p. 41). ‘Despite the higher total sum of individual goodness in h_2 , it is possible that someone impressed by the vision of total equality in h_1 will swear by Babeuf that it is superior to h_2 ’. But in considering the choice between h_3 and h_4 , ‘all that vision of equality is now gone, and in this rather mundane choice it is possible that the same man may not really be able to say that he prefers h_3 to h_4 ’. But if Sen’s character judges that h_1 is better than h_2 , and also that h_3 is not better than h_4 , he is forced to deny strong separability. I will call such judgements *Allais-like*, and the claim that they somehow show that egalitarianism and strong separability are incompatible *the argument from Allais-like cases*.

However, we will soon see that the argument from Allais-like cases is fairly weak, so it is worth looking at another argument for the

incompatibility of egalitarianism with strong separability. Suppose members of some subpopulation are each unaffected by the choice between two histories. Strong separability implies that we can ignore this subpopulation in ranking the two histories. But egalitarians care about the patterns of equality and inequality, and those patterns are global. So how good the two histories are for the members of the subpopulation could make a difference to the extent of inequality in the two histories, in one case for the better and in another case for the worse. Hence egalitarians have to deny strong separability. I will call this the *argument from patterns*. I believe it is this argument which leads many people to think that egalitarianism has to reject strong separability.

Thus we have two rationales for denying the compatibility of egalitarianism and strong separability: the argument from Allais-like cases, and the argument from patterns. Since it is very far from obvious what theory E implies about strong separability, these rationales generate a worry about (E).

12. Criticisms I

I think Sen's discussion of Allais-like cases was only intended to show that egalitarians can reject strong separability, and in that respect it seems convincing. But as Broome (forthcoming) suggested, it is hard to see a reason why someone concerned with equality has to make an Allais-like judgement about any particular Allais-like case.

Any Allais-like case involves four histories, k_1 through k_4 , meeting the conditions needed for the definition of strong separability to apply, with two further properties: (i) k_1 is better in terms of minimal equality than k_2 , while k_3 is not better than k_4 in terms of minimal equality; and (ii) k_2 has a greater sum of individual goodness than k_1 , and k_4 has a greater sum than k_3 .

However, if egalitarians are characterized as having two concerns, increasing the sum of individual goodness and reducing inequality, someone does not fail to be an egalitarian merely by caring a lot about the sum and relatively little about inequality. So it is obvious that in any particular Allais-like case, some egalitarian could think that because of their greater sums, k_2 is better than k_1 , and k_4 is better than k_3 , despite whatever advantages in terms of equality k_1 and k_3 enjoy. Such an egalitarian therefore does not hold an Allais-like judgement about the particular case in question.

It should now be obvious that there is little prospect for defending the following claim: there exists an Allais-like case about which every form of egalitarianism has to make an Allais-like judgement. So someone who wishes to appeal to Allais-like cases to argue for the incompatibility of egalitarianism and strong separability is going to have to retreat to defending the logically weaker claim: every form of egalitarianism has to make an Allais-like judgement about some Allais-like case. However, this weaker claim is not much different from the blank denial of the compatibility of egalitarianism and strong separability, and it can only preach to the converted. Thus appeals to Allais-like cases do not take us far in determining whether egalitarianism is incompatible with strong separability.

Perhaps someone who wishes to defend the incompatibility of egalitarianism with strong separability should instead rest her case on the argument from patterns. An advantage of this move is that we no longer need to work within the framework of the risk-free interpretation of the fundamental egalitarian idea. Unfortunately, however, despite the fact that something along the lines of the argument from patterns has often been endorsed (including, unfortunately, in McCarthy 2008), it is a bad argument.

Whether or not it is bad, however, it is a bit vague, so we should begin by sharpening it. In talking about the presence of some group making a difference to the overall patterns of inequality, the argument from patterns is obviously talking about a formal property of the relation of one history being at least as good in terms of equality as another. More precisely, since it is making a claim about all forms of egalitarianism, we should interpret it as making a claim about the minimal betterness in terms of equality relation, which I have claimed is \succeq_e . Now one thing the Allais-like example does show is that \succeq_e is not strongly separable. For while $h_1 \succ_e h_2$, it is not the case that $h_3 \succ_e h_4$. So it is natural to interpret the argument from patterns as claiming, correctly, that the minimal betterness in terms of equality relation is not strongly separable, and from that inferring that according to egalitarianism, the risk-free betterness relation is not strongly separable.

However, unless most egalitarians are to be convicted of a simple inconsistency, this last inference is dubious. To explain, whereas strong separability across individual outcomes says that we can ignore unaffected groups of any size, what I will call weak separability across individual outcomes says that we can ignore unaffected groups provided they contain all members of the population but one. So as

the terminology suggests, strong separability entails weak separability, but not vice versa. More precisely,

Definition 8

Suppose histories h_1 through h_4 are such that for some individual i , (i) h_1 is exactly as good for i as h_3 , and h_2 is exactly as good for i as h_4 , and (ii) for every individual j apart from i , h_1 is exactly as good for j as h_2 , and h_3 is exactly as good for j as h_4 . Then a preorder \succeq on histories is *weakly separable across individual outcomes* iff in all such cases: $h_1 \succeq h_2$ iff $h_3 \succeq h_4$.

Now not only is \succeq_e not strongly separable, it is not even weakly separable, as an example which will be important later on illustrates.

Example 1

Let $h_1 = [1, 1]$, $h_2 = [0, 1]$, $h_3 = [1, 0]$ and $h_4 = [0, 0]$. Then letting individual i be person 1, we see that the definition of weak separability applies to h_1 through h_4 . But $h_1 \succ_e h_2$ while $h_4 \succ_e h_3$. Thus \succeq_e is not weakly separable.

Should egalitarians conclude that because \succeq_e is not weakly separable, neither is the risk-free betterness relation? Not if they want to be moderate, at least in a slightly spruced-up sense. For suppose that everyone apart from i is unaffected by the choice between h_1 and h_2 . Moderacy implies that if h_1 is better/exactly as good as/worse than h_2 for i , then h_1 is better/exactly as good as/worse than h_2 . To make room for incompleteness, it is very natural to add: if h_1 is incomparable with h_2 in terms of betterness for i , then h_1 is incomparable with h_2 in terms of betterness. Then on this spruced-up version of moderacy, which is how I will understand it in what follows, moderacy immediately implies that the risk-free betterness relation is weakly separable. For if everyone apart from i is unaffected by the choice between h_1 and h_2 , then according to moderate egalitarianism, h_1 is at least as good as h_2 iff h_1 is at least as good for i as h_2 . This strongly suggests that one cannot straightforwardly claim that according to egalitarianism, the risk-free betterness relation inherits the separability properties of the minimal betterness in terms of equality relation. But that is just what the argument from patterns naively seems to claim. We will return to this topic in section 15, but now it is time to look at some arguments in the other direction.

13. Compatibility with strong separability?

Some writers claim that egalitarianism is compatible with strong separability (see, for example, Atkinson 1970, Tungodden 2003). This

claim is generally made without talking about risk, and in fact there is a simple argument for compatibility if we adopt the risk-free interpretation of the fundamental egalitarian idea.

Like the risk-involving interpretation, the risk-free interpretation features the minimal betterness in terms of equality relation. Since we identified that relation with \succeq_e for the risk-involving interpretation, we should surely do the same for the risk-free interpretation. But let g be an individual goodness measure. Then the risk-free interpretation implies that the following claim is a version of egalitarianism.⁵

(W) The risk-free betterness relation is represented by $w \circ g_1 + \dots + w \circ g_n$ for some strictly increasing and strictly concave function w

Proof: given the identity of \succeq_e with the minimal betterness in terms of equality relation, (W) obviously entails (a) through (c) of the risk-free interpretation.

However, if the risk-free betterness relation is represented by $w \circ g_1 + \dots + w \circ g_n$ for some function w , it is strongly separable. Hence the risk-free interpretation implies that egalitarianism is compatible with strong separability. I will call this *the argument from the risk-free interpretation*.

14. Criticisms II

The argument from the risk-free interpretation is valid, but one of its intermediate conclusions is troubling. For working within the same kind of risk-free framework, the priority view has often been defined as the view which claims that (W) is true (for discussion of such definitions, see McCarthy 2013). And although Parfit (2000) never offered a formal definition of the priority view, he repeatedly contrasts it with egalitarianism, claiming, for example, that unlike egalitarianism, the priority view is not concerned with patterns between people and is not subject to the levelling-down objection. So it should at least be surprising that a popular definition of the priority view makes it a form of egalitarianism (Jensen 2003).

Let us probe this further. The informal claim which seems to point towards the most fundamental distinction between egalitarianism and

⁵ A function $f: \mathbb{R} \rightarrow \mathbb{R}$ is *strictly concave* just in case for any $x, y \in \mathbb{R}$, and any $\alpha \in (0, 1)$, $f(\alpha x + (1 - \alpha)y) > \alpha f(x) + (1 - \alpha)f(y)$. Recall that $w \circ g_1 + \dots + w \circ g_n$ is a function on histories whose value on any history h is $w(g_1(h)) + \dots + w(g_n(h))$.

the priority view is the egalitarian idea that equality is always in one respect better than inequality. It was this idea which motivated the argument from patterns, and which underlay Parfit's levelling-down objection to egalitarianism. By contrast, prioritarists reject this idea. For example, they think that there is no respect in which $[0, 0]$ is better than $[1, 0]$. As the same example illustrates, prioritarists will also reject the idea that downwards contractions are always improvements in one respect. But to save words, I will focus on equality.

It is incumbent upon any adequate interpretation of egalitarianism, and of the contrast between egalitarianism and the priority view, to find a way of adequately expressing the idea that for egalitarians, equality is always in one respect better than inequality. And it may seem that the risk-free interpretation of the fundamental egalitarian idea succeeds in doing this. For, roughly speaking, the risk-free interpretation seems to imply, among other things, that for any histories h_1 and h_2 , (a) if h_1 contains perfect equality and h_2 contains inequality, then h_1 is better in terms of equality than h_2 , and (b) h_1 being better in terms of equality than h_2 counts towards h_1 being better than h_2 .

However, this is an illusion. The risk-free interpretation does not do justice to the idea that according to egalitarians, equality is in one respect better than inequality. For it is obvious that once we accept that idea, we should claim that in Myerson's example L_E is better than L_F . But it is consistent with the risk-free interpretation that L_E and L_F are equally good, even if we accept the usual assumption to do with risk, namely that the betterness relation satisfies the expected utility axioms.

To demonstrate this, we already know that the risk-free interpretation implies that (W) is a form of egalitarianism. Adding the claim that the betterness relation satisfies the expected utility axioms to (W) is consistent with the claim that the betterness relation is represented by taking the expected value of $w \circ g_1 + \dots + w \circ g_n$. But no matter how individual goodness measures are defined, that claim implies that L_E and L_F are equally good.

We started with an argument to the effect that egalitarianism is compatible with strong separability because the risk-free interpretation of the fundamental egalitarian idea implies that it is. But rather than conclude that this shows that egalitarianism and strong separability are compatible, we might rather conclude this. Because it does not genuinely capture the idea that equality is in one respect better than inequality, the idea which seems to be the most critical for understanding the way in which egalitarianism is concerned with

patterns, the risk-free interpretation is an unreliable guide to what egalitarianism implies about strong separability.

15. Interim conclusions about patterns

I will begin with some lessons from the failure of the arguments for and against the compatibility of egalitarianism with strong separability. I will focus on the argument from patterns and the argument from the risk-free characterization because they make complementary mistakes: they each neglect one side of the egalitarian coin.

The argument from patterns pays a lot of attention to equality. In effect, it assumes that egalitarian risk-free betterness relations are similar enough to the minimal betterness in terms of equality relation for us to be able to claim that the former inherits the separability properties of the latter. Now it seems right to try to somehow appeal to the separability properties of the minimal betterness in terms of equality relation to understand what egalitarians have to say about the separability properties of the betterness relation. But the argument from patterns does this too crudely. Because egalitarianism cares about well-being as well as equality, it is hard to see why the risk-free betterness relation should automatically inherit the separability properties of a preorder which is solely concerned with equality. Thus the argument from patterns neglects the fact that egalitarianism values well-being.

The argument from the risk-free interpretation seems to improve on this. It pays attention to well-being, and appears to pay attention to equality as well. But the most basic aspect of an egalitarian concern with equality is the view that equality is always in one way better than inequality. However, we saw that although the risk-free interpretation appears to capture this idea, that appearance is illusory. Thus the argument from the risk-free interpretation does not succeed in capturing the way egalitarianism values equality.

We began the discussion of patterns with a challenge to (E). There are strongly held views about what egalitarianism implies about patterns between people, and it is not obvious whether theory E coheres with these views, putting (E) in doubt. But looking at the argument from patterns and the argument from the risk-free interpretation together, we find tacit support for something which should now seem obvious. Egalitarianism is concerned with two values, well-being and equality. So any conclusion about the compatibility of egalitarianism with this or that view about patterns should emerge from an account of egalitarianism which fully captures its concern with these values.

This, I have claimed, is provided by the characterization of the core of egalitarianism provided earlier, the conjunction of the equality principle and the equality-neutral principle of personal good. So say not that theory E should be judged by whether it fits with some independent view about what egalitarians should say about patterns. Say rather that what egalitarians should say about patterns is determined by what theory E does say. We now turn to that.

16. Weak separability across individual lotteries

In both lotteries in Myerson's example, each person faces an individual lottery of the form $[\frac{1}{2}, 1; \frac{1}{2}, 0]$. But egalitarians will not be indifferent to how the two individual lotteries are combined, for they think that L_E is better than L_F . Thus for egalitarians, some ways of combining fixed individual lotteries are better than others, resulting in a failure of a certain kind of separability.

A natural way of trying to characterize the form of separability which egalitarians have to reject is as follows: they have to deny that the contribution an individual lottery makes to overall goodness is independent of context, where the context is given by fixing the combination of individual lotteries other individuals face.

More precisely, say that two lotteries L_1 and L_2 are *equivalent relative to a group of individuals* just in case they have the form $L_1 = [p_1, h_1; \dots; p_m, h_m]$ and $L_2 = [p_1, k_1; \dots; p_m, k_m]$, and for each individual i in the group and for all $j = 1, \dots, m$, h_j is exactly as good for i as k_j .

Definition 9

Suppose lotteries L_1 through L_4 are such that for some individual i , (a) L_1 and L_3 are equivalent relative to i , as are L_2 and L_4 ; and (b) L_1 and L_2 are equivalent relative to the group containing everyone apart from i , as are L_3 and L_4 . Then a preorder \succsim on lotteries is *weakly separable across individual lotteries* iff in all such cases: $L_1 \succsim L_2$ if and only if $L_3 \succsim L_4$.

Myerson's example can now immediately be used to show that egalitarians have to deny that the betterness relation is weakly separable across individual lotteries.⁶

⁶ The only discussion I know of this form of separability is in Broome 1989, under the name of 'weak prospect separability'. But Broome's discussion is brief, and he regards the rejection of strong separability across individual outcomes as the defining feature of egalitarianism, a view maintained in Broome forthcoming.

Example 2

Let $L_1 = [\frac{1}{2}, [1, 1]; \frac{1}{2}, [0, 0]]$, $L_2 = [\frac{1}{2}, [0, 1]; \frac{1}{2}, [1, 0]]$, $L_3 = [\frac{1}{2}, [1, 0]; \frac{1}{2}, [0, 1]]$, $L_4 = [\frac{1}{2}, [0, 0]; \frac{1}{2}, [1, 1]]$. Then letting person 1 be i in the preceding definition, if the betterness relation is weakly separable across individual lotteries, then: L_1 is at least as good as L_2 if and only if L_3 is at least as good as L_4 . But for egalitarians, L_1 is better than L_2 while L_4 is better than L_3 , as implied by the equality principle.

Thus I suggest the following:

Proposal 4

The fundamental form of separability rejected by egalitarianism is that the betterness relation is weakly separable across individual lotteries

The basic idea is that because they care about certain patterns, even when it is fixed which individual lottery each person will face, egalitarians will prefer some ways of combining these individual lotteries over others. In particular, they will prefer combinations which will improve the patterns in each history which could result. This forces egalitarians to reject weak separability across individual lotteries.

Ideas to do with risk often have several variants, and Proposal 4 is no exception. But ignoring these, can we be so sure that it is the rejection of weak separability across individual lotteries which lies at the heart of egalitarianism, rather than the rejection of some rather different form of separability? In defence of Proposal 4 I will offer three replies.

First, I find it hard to see what this other form of separability would be. But even if there is a decent candidate, it seems unlikely that its rejection could be derived as quickly and directly from the core of egalitarianism as the rejection of weak separability across individual lotteries. Second, section 18 will argue that Proposal 4 forms part of an appealing story about how egalitarianism relates to other distributive views, and that surely counts in its favour. Third, Proposal 4 is in some sense just what we should expect.

To elaborate, recall that the argument from patterns tried to derive separability properties of egalitarian betterness relations from separability properties of betterness in terms of equality. That is a natural and obvious strategy, and it was just that in neglecting the egalitarian concern with well-being, the argument from patterns executed it too crudely. By contrast, the argument for Proposal 4 handles the existence of that concern by appealing exclusively to lotteries which are all

equally good in terms of well-being, so that there is nothing to choose between them in terms of well-being. This also corrects the mistake of the argument from Allais-like cases, which fails because it allows us to choose between the histories concerned on the grounds of well-being. Nevertheless, the argument for Proposal 4 takes its cue from the argument from patterns in providing a direct and simple derivation of separability properties of egalitarian betterness relations from separability properties of betterness in terms of equality. Thus Example 1, used to show that the minimal betterness in terms of equality relation is not weakly separable across individual outcomes, immediately converts to Example 2, used to show that egalitarian betterness relations are not weakly separable across individual lotteries. Thus Proposal 4 fills the role we should have anticipated.

17. Strong separability across individual outcomes

We can now return to the question of the compatibility of strong separability with egalitarianism. Although the answer will feed into later philosophical arguments, the question is by now a purely mathematical one, and is just whether theory E is consistent with strong separability across individual outcomes.

While working in a risk-free framework, some writers have claimed that leximin is a form of egalitarianism in order to argue that egalitarianism is compatible with strong separability (see, for example, Tungodden 2003).⁷ Leximin is indeed a form of egalitarianism if the risk-free interpretation of the fundamental egalitarian idea is correct, but the worries about the appeal to the risk-free interpretation stated in section 14 remain. In addition, there is the further worry that although moderate, leximin is nevertheless rather an extreme view. Even if one were to grant that leximin is a form of egalitarianism, this would not by itself establish that less extreme egalitarian views which almost all egalitarians find more plausible are compatible with strong separability across individual outcomes. It is typically assumed that the way to characterize moderate views which are less extreme than egalitarianism is to require that the

⁷ Leximin gives absolute priority to benefiting the worst off; then if that is no longer possible, the next worse off, etc. It is easily verified that it is strongly separable across individual outcomes.

risk-free betterness relation is continuous,⁸ so we will investigate theory E with that additional assumption.

When Broome (1989) discussed weak separability across individual outcomes, he said he would be ‘surprised’ if egalitarians could reject weak separability across individual lotteries but accept strong separability across individual outcomes. Broome’s intuition is very natural, and it is remarkable that it turns out to be wrong.

Call *strongly separable continuous theory E* the version of theory E which is continuous and accepts strong separability across individual outcomes. Recall that we are assuming a population of individuals $1, \dots, n$. To skip a technicality, I will assume at least three people. For any strictly increasing and strictly concave function w , let E_w be the function which maps lotteries to real numbers defined as follows. For any lottery $L = [p_1, h_1; \dots; p_m, h_m]$,

$$E_w(L) := \sum_j p_j w^{-1} \left(\frac{1}{n} [w(u_1(h_j)) + \dots + w(u_n(h_j))] \right)$$

Then slightly reducing generality for readability, we have the following theorem:

Theorem 1

Assume a fixed population $1, \dots, n$, $n \geq 3$. Assume also that the extended individual betterness relation satisfies the expected utility axioms, so that it can be represented by some expectational function U , and suppose the set of utility profiles is the set of n -tuples of real numbers. Then: strongly separable continuous theory E \Rightarrow for some strictly increasing and strictly concave function w , E_w represents the betterness relation.

This is proved in McCarthy and Thomas MS.⁹ Theorem 1 shows that on the assumption that egalitarianism is theory E, egalitarianism is consistent with strong separability across individual outcomes. It immediately implies the following:¹⁰

⁸ Continuity will be formally defined in almost any textbook covering utility theory, such as Ok 2007. Informally, the idea is that if a (history with) utility profile u is better than a utility profile v , then any utility profile which is sufficiently close to u will also be better than v , and any profile which is sufficiently close to v will be worse than u . Closeness is understood in terms of the usual way of measuring distance in Euclidean spaces.

⁹ The same result, but with theory W (defined in Sect. 9.2) substituted for theory E, was independently presented by the author and by Marc Fleurbaey in 2007, later published in Fleurbaey 2010. But as already noted, theory W is too weak to count as egalitarianism.

¹⁰ By Lemma 1 and the fact that since w is strictly increasing, so too is w^{-1} .

Corollary 1

Assume the conditions of Theorem 1. Then strongly separable continuous theory E implies that for some strictly increasing and strictly concave function w , the risk-free betterness relation is represented by $w \circ u_1 + \dots + w \circ u_n$. Conversely, for any strictly increasing and strictly concave function w , the risk-free betterness relation is represented by $w \circ u_1 + \dots + w \circ u_n$ according to some version of strongly separable continuous theory E.

Given (E), this corollary teaches us that strongly separable continuous egalitarianism is moderate. But beyond that, it comes in a wide variety of strengths. For without going into details (see Bosmans 2007 for an overview), orderings represented by $w \circ u_1 + \dots + w \circ u_n$ for some strictly increasing and strictly concave w range from arbitrarily close to utilitarianism to arbitrarily close to leximin.

18. Taxonomy and risk

This section has two parts. The first asks what the contrast between egalitarianism and other distributive views looks like if we accept (E). It says more about patterns, and Proposal 4 in particular. The second addresses a natural objection to (E). As already noted, many discussions of egalitarianism never mention risk. So it may seem implausible that an account of egalitarianism should feature risk essentially. But drawing on the first part of the discussion, I will argue that talking about risk in our taxonomy of basic distributive views costs us nothing and benefits us plenty.

To do all this I need to discuss not only egalitarianism, but also utilitarianism, the priority view, and fairness. Each of these topics deserves its own article, so the presentation is going to have to be condensed. To simplify I will continue with the bells and whistles built into the statement of the utilitarian and egalitarian premisses, and in particular assume that any n -tuple of real numbers is a utility profile. I will also ignore hybrids, such as views which are concerned with both equality and priority. So when I refer to a concern with equality, for example, I mean something like a departure from utilitarianism solely in the direction of caring about equality. To save words, I will often say 'weak separability' for 'weak separability across individual lotteries.'

18.1 Utilitarianism and the priority view

Given (E), we know where utilitarianism disagrees with egalitarianism: utilitarianism rejects the equality principle. But we can now say more.

Egalitarians reject weak separability because of a particular concern with patterns between people, exemplified by the view that equality is always in one respect better than inequality. But other concerns with patterns would lead to the rejection of weak separability. For example, inequality will reject weak separability because they think that inequality is always in one respect better than equality. So in Myerson's example they will think L_F is better than L_E . However, it is very hard to see any rationale for rejecting weak separability aside from being concerned with patterns. We are accustomed to thinking that utilitarianism is not concerned with patterns, so we should expect the acceptance of weak separability to be close to the surface of the utilitarian premisses. And indeed it is, for anonymous Pareto immediately implies weak separability.

Let us turn now to the priority view. McCarthy (2013) argues that it is not possible to make satisfactory sense of the priority view in a risk-free framework, with particular focus on the role of evaluative measurements. So here I will look at how to understand the way the priority view contrasts with egalitarianism and utilitarianism in a risk-involving framework.

The priority view has been said to contrast with egalitarianism in two ways (Sect. 14): unlike egalitarianism, the priority view is neither concerned with patterns nor subject to a levelling-down objection. The first contrast has often been analysed in terms of strong separability across individual outcomes: the priority view accepts it, while egalitarianism rejects it. But given (E), we can now see this gets it wrong about egalitarianism. A better analysis arises from saying that like utilitarianism, the priority view accepts weak separability across individual lotteries.¹¹

But what of the second contrast? In section 20 I will claim that because of its rejection of weak separability, egalitarianism is subject to a serious levelling-down objection, one which is stronger than Parfit's original objection. Thus borrowing the conclusion of that section, Proposal 4 forms part of an appealingly simple view about how egalitarianism contrasts both with utilitarianism and the priority view.

I need to sketch what I take to be the contrast between the priority view and utilitarianism. Consider a world containing only one person, Alice. Suppose she has to choose between a history $h_p = [1]$ and a

¹¹ I think Broome (forthcoming) was the first to suggest using Myerson's example to distinguish between egalitarianism and the priority view. But rather than conclude that it is weak separability across individual lotteries which separates them, he concludes, mistakenly, that it is strong separability across individual outcomes.

lottery $L_P = [\frac{1}{2}, 2; \frac{1}{2}, 0]$. I take the essence of the priority view to be exemplified by the claim that even though h_P and L_P are equally good for Alice, h_P is nevertheless better than L_P .¹² I call the principle which generalizes this claim the *priority principle* (McCarthy forthcoming c). I do not have the space to spell out the details, so I am going to have to treat this term as a placeholder for an idea. But regardless of the details, the priority principle is going to conflict with utilitarianism by being inconsistent with anonymous Pareto.

18.2 Equality of prospects

Consider an example due to Diamond (1967), with L_F repeated for convenience.

L_F	<i>heads</i>	<i>tails</i>	L_U	<i>heads</i>	<i>tails</i>
<i>Alice</i>	1	0	<i>Alice</i>	1	1
<i>Bob</i>	0	1	<i>Bob</i>	0	0

There is a case for saying that although L_F and L_U are bound to produce equally good outcomes, L_F is better than L_U . For L_F gives Alice and Bob equal chances, while L_U guarantees that Alice will win and Bob will lose. Or to put it another way, Alice and Bob face equal prospects under L_F , but unequal prospects under L_U .

Many people think that just as one can value equality in the distribution of goods, one can also value fairness in the distribution of goods. I have been claiming that Myerson's example embodies what it is to value equality, and if I were to hold the view just mentioned about fairness, I would claim that Diamond's example embodies what it is to value fairness. Indeed, Diamond's original idea was that L_F is better than L_U because it is fairer, and Diamond's example lies at the heart of Broome's rich account of distributive fairness (Broome 1990, 1991).

Unfortunately, I no longer hold the background view about fairness. It is not the case that one can value distributive fairness in much the same way that one can value distributive equality. I will say more about this in section 21.2. Instead, I prefer to say that Myerson's example embodies what it is to value equality of outcome, and

¹² That the priority view has this sort of consequence has been argued in Rabinowicz 2001, 2002, McCarthy 2008, and Otsuka and Voorhoeve 2009. But what these works see as the consequence of some independent characterization of the priority view, I see as the essence.

Diamond's example embodies what it is to value equality of prospects.¹³

There are clearly some fascinating parallels between Myerson's and Diamond's examples. But no matter what terminology one uses to describe these examples, the concerns they embody are independent. Suppose, for example, we say that Myerson's example is about equality, and Diamond's is about fairness. Then there is nothing obviously incoherent about being concerned with equality but not fairness, and nothing obviously incoherent about being concerned with fairness but not equality. Diamond's and Myerson's examples are concerned with different but equally fundamental ideas.

Now although I would prefer to say that Myerson's example is about equality of outcome, and Diamond's is about equality of prospects, it will be easier to compare my views with the literature if instead I henceforth say that Myerson's example is about equality and Diamond's is about fairness. This will be harmless, but we will return to the issue.

Just as the equality principle emerged as a natural generalization of the idea that in Myerson's example, L_E is better than L_F , we can generalize the idea that in Diamond's example, L_F is better than L_U .¹⁴ I will call the hypothesized result the *fairness principle*. But as with the priority principle, I am going to have to leave this idea as a placeholder.

I have claimed that theory E departs from utilitarianism solely in the direction of a concern with equality, and that a concern with fairness is independent of a concern with equality. If so, there must be something about theory E which expresses indifference to fairness. But what? The answer was supplied by Diamond (1967), who noted that (*pace* Broome (1991)), thinking L_F is better than L_U requires rejecting the idea that the betterness relation satisfies the central axiom of expected utility theory, the strong independence axiom. So it is in accepting that betterness satisfies strong independence that theory E expresses indifference to fairness.

Section 16 argued that Proposal 4 leads to an attractive internal picture of egalitarianism. Nothing in Diamond's example motivates rejecting weak separability, so weak separability is going to be accepted

¹³ To a reasonable approximation, the 'equality of outcome versus equality of prospects' distinction is also commonly marked as the 'ex post equality versus ex ante equality' distinction.

¹⁴ How to generalize it is discussed in McCarthy and Thomas MS.

not just by utilitarianism and prioritarianism, but also by concerns with fairness. Thus Proposal 4 is part of an appealingly simple picture of where egalitarianism contrasts with other distributive views.

18.3 A case study

We can now turn to the objection that it is implausible that an account of the core of egalitarianism, or indeed of any distributive view, should essentially involve risk. To respond to this objection I am going to contrast taxonomical life with risk and without it. However, I cannot hope to do this in any detail. Instead, I am going to focus on one family of risk-free taxonomical views, and compare it with what I hope will seem to be a natural development of the risk-involving ideas just sketched.

18.3.1 A risk-free taxonomy

Let \mathcal{W} be the set of orderings of histories represented by $w \circ u_1 + \dots + w \circ u_n$ for some strictly increasing and strictly concave function w . Many philosophers discussing the ethics of distribution tacitly or explicitly take the individual utility functions u_1, \dots, u_n to be individual goodness measures. Then without saying any more about risk (for to refer to the utility functions is already to say something about risk), the following views are reasonably popular. First, the set of egalitarian orderings of histories includes \mathcal{W} . This follows from the risk-free interpretation of the fundamental egalitarian idea. Second, the set of orderings of histories concerned with fairness also includes \mathcal{W} . For example, this naturally models Broome's view that fairness requires distributing goods in proportion to how much each person would benefit them (Broome 1990, 1991; see also Segal 2006). Third, the set of prioritarian orderings of histories is identical to \mathcal{W} ; see McCarthy 2013 for references.

In symbols, let \mathcal{E}_h , \mathcal{F}_h and \mathcal{P}_h be the sets of orderings of histories which can arise from concerns with equality, fairness, or priority respectively. Then the claims just made are

$$(2) \quad \mathcal{W} \subseteq \mathcal{E}_h, \quad \mathcal{W} \subseteq \mathcal{F}_h, \quad \text{and} \quad \mathcal{W} = \mathcal{P}_h$$

Thus any ordering in \mathcal{W} can be interpreted as concerned with equality, fairness, or priority. In fact, it could also be interpreted as a hybrid of any two or all three of those concerns, but I am ignoring hybrids.

18.3.2 A risk-involving taxonomy

I am confident that the understanding of concerns with fairness and priority given in 18.1 and 18.2 can be axiomatized along the same lines of the axiomatization of egalitarianism given in this article (no use of individual goodness measures, using only comparatives, dispensability of simplifying assumptions, etc.). But I have only had the space to gesture towards accounts of the two most distinctive features of those concerns, the fairness principle and the priority principle. So I am in no position to provide the details of these axiomatizations. Instead, I am just going to elaborate on three consequences of the resulting risk-involving taxonomy.

The first is that orderings, now of lotteries, which express concerns with equality, fairness, or priority are entirely distinct. In other words, no ordering can be interpreted as both being concerned with equality and with fairness, for example. This follows from the differing ways the three concerns necessarily treat the examples L_E , L_F , L_U , L_P , and h_P . In symbols, let \mathcal{E} , \mathcal{F} , and \mathcal{P} be the sets of orderings of lotteries which can arise from concerns with equality, fairness, or priority respectively. Then the first consequence of the risk-involving taxonomy is

$$(3) \quad \mathcal{E} \cap \mathcal{F} = \emptyset, \quad \mathcal{E} \cap \mathcal{P} = \emptyset, \quad \text{and} \quad \mathcal{F} \cap \mathcal{P} = \emptyset$$

The second is that concerns with fairness and priority have natural representations. For any strictly increasing and strictly concave function w , let F_w and P_w be functions which map lotteries to real numbers defined as follows. For any lottery $L = [p_1, h_1; \dots; p_m, h_m]$, $F_w(L) := \sum_i w(U_i(L))$ and $P_w(L) := \sum_j \sum_i p_j w(u_i(h_j))$. Let \mathcal{E}_w , \mathcal{F}_w and \mathcal{P}_w be the sets of orderings of lotteries represented by E_w , F_w and P_w for some strictly increasing and strictly concave w . Then the risk-involving taxonomy yields¹⁵

$$(4) \quad \mathcal{E}_w \subseteq \mathcal{E}, \quad \mathcal{F}_w \subseteq \mathcal{F}, \quad \text{and} \quad \mathcal{P}_w = \mathcal{P}$$

The third is that despite their essential disagreement about lotteries, when orderings concerned with equality, fairness, and priority are restricted to histories, there is a lot of overlap. In particular, (4) immediately implies (2).¹⁶

¹⁵ The first of the following claims has been argued for here, the second is a consequence of McCarthy and Thomas MS, and the third is argued for in McCarthy forthcoming c.

¹⁶ Because for any fixed strictly increasing and strictly concave w , the functions E_w , F_w , and P_w represent exactly the same ordering when restricted to histories.

18.3.3 Comparison

Suppose we accept the risk-free taxonomy, thinking we do not need to talk about risk to make sense of basic distributive ideas. What are we to make of the fact that any ordering in \mathcal{W} can be interpreted as expressing a concern with equality, a concern with fairness, and a concern with priority?

There are a number of risk-free views about the relationship between these concerns which could be seen as partial responses. For example, some writers identify concerns with equality and concerns with fairness. Others claim that concerns with priority are not a genuine alternative to concerns with equality. Still others conclude that the distinction between the two should be made at the level of justification.

None of these is an entirely happy position. Concerns with equality, fairness, and priority at least sound as if they are distinct. But drawing distinctions at the level of justification while denying that the distinctions generate disagreements at the concrete level of cases raises a worry that the distinctions are not real.

Matters are different, however, if we accept the risk-involving taxonomy. We have seen that orderings of lotteries can be interpreted as being concerned with at most one of equality, fairness, and priority. So the risk-involving taxonomy enables us to cash out the difference between these concerns at the level of concrete disagreement about cases in a way which clearly relates to their differing justifications. Moreover, it also explains why, at the risk-free level, there is so much overlap. This surely makes the risk-involving taxonomy superior. Exactly the same point would have emerged had we been talking about equality of outcome and equality of prospects instead of equality and fairness. The risk-involving taxonomy would distinguish concerns with equality of outcome, equality of prospects, and giving priority to the worse-off at the level of risk while explaining their substantial risk-free overlap.

A defender of the risk-free taxonomy might object. She has to concede that the risk-involving taxonomy brings benefits. But she may also claim that it comes with outweighing costs: a significant increase in complexity, and the *prima facie* unnaturalness of the idea that many basic distributive views are essentially about risk. In response, whether the apparent unnaturalness is anything more than a surface appearance is up for grabs. But, more important, everyone eventually has to face the complexities to do with risk, so it is a false economy not to

help ourselves from the outset to the expressive resources talk about risk provides.

19. Individual goodness measures

Recall from section 2 that individual goodness measures are functions on individual/history pairs which somehow purport to measure how good the history is for the individual in question. I have claimed that egalitarianism can and should be defined in a risk-involving framework without taking the existence of individual goodness measures for granted. But there remains the question of what the full story about individual goodness measures should be, and how the risk-free interpretation of the fundamental egalitarian idea relates to the risk-involving interpretation.

The risk-free interpretation accords with a popular approach to the ethics of distribution. Moral philosophers often take it for granted that it makes sense to talk about units of individual goodness, and that various distributive principles which exploit this talk are then reasonably transparent to intuition. Call this the *quantitative approach*. For example, writers who take the quantitative approach think it is intuitively clear that if equality is valuable, an equal distribution is better than an unequal distribution of the same sum of individual goodness.

I disagree with the quantitative approach. In my view, distributive theories should first be formulated in terms of evaluative comparatives like the betterness relation and individual betterness relations without appealing to evaluative measures like individual goodness measures. Evaluative measures then emerge from the distributive theories in virtue of the role they play in the best quantitative description of those theories. Call this the *comparative approach*. My defence of (E) fits with this approach.

Although the comparative approach is nonstandard for moral philosophy, it fits well with the way nonevaluative measures are treated in the physical and social sciences (see, for example, Krantz et al. 1971, Narens 2007, Roberts 2009). I argue for the comparative approach elsewhere (McCarthy forthcoming b), but I hope that its fit with the literature on measurement makes it worth exploring against the background of (E).

The crucial question, on the comparative approach, is what makes some quantitative description of distributive theories the best description. In my view, there are two desiderata: simplicity, and fit with

established discourse. Our discussion of egalitarianism will illustrate the latter.

Borrowing a term from Broome 1991, let *Bernoulli's hypothesis* be the claim that individual goodness measures are positive affine transformations of u . Bernoulli's hypothesis makes one lottery better for an individual than another just in case it gives her a greater expected amount of individual goodness. Consider the following:¹⁷

Theorem 2

Suppose the set of utility profiles is the set of n -tuples of real numbers. Then theory E \Rightarrow [the risk-free interpretation of the fundamental egalitarian idea \Leftrightarrow Bernoulli's hypothesis].

Let us summarize the dialectical situation. For the reasons given in section 3, the only obvious constraint on individual goodness measures is that they have to represent the risk-free extended individual betterness relation. This only makes individual goodness measures unique up to increasing transformation. But as already noted, if we want to talk about units of individual goodness, individual goodness measures need to be unique up to positive affine transformation. So on the comparative approach we face a choice about how to define them. But given that the risk-free interpretation of the fundamental egalitarian idea encapsulates widely held ideas about what egalitarianism is, we obviously do better to exercise that choice in a way which makes the risk-free interpretation come out true. Given (E), Theorem 2 therefore tells us to adopt Bernoulli's hypothesis.

I said there were two desiderata for choice of individual goodness measures: simplicity, and fit with established discourse. I have argued that at least when we consider egalitarianism, fit with established discourse leads to Bernoulli's hypothesis. Several writers have in effect argued that we get a simpler quantitative account of the ethics of distribution by adopting Bernoulli's hypothesis rather than some competing measure of individual goodness (Broome 1991, Hammond 1991, McCarthy 2007). So happily, the two desiderata point in the same direction.

To some extent, then, the risk-free interpretation of the fundamental egalitarian idea turns out to be correct. But there are important caveats to this claim. First, the adoption of Bernoulli's hypothesis does not make the risk-free interpretation entail that in Myerson's example, L_E is better than L_F . So the risk-free interpretation still does not

¹⁷ A proof is available from the author.

capture the most central egalitarian idea. Second, the claims about egalitarianism the risk-free interpretation makes are derived, not definitional. Third, the adoption of Bernoulli's hypothesis presupposes that the extended individual betterness relation satisfies the expected utility axioms. But it is very unclear how to make sense of the risk-free interpretation in the absence of that presupposition. By contrast, although there is more work to do, I have shown how my account of the core of egalitarianism can do without this presupposition. Given how widespread quantitative talk about individual goodness is in discussions of egalitarianism and other distributive views, it is probably somewhat useful to make sense of it. But far more important, in my view, is the fact that we can perfectly well theorize about egalitarianism and other views without needing it.

20. Strong levelling down

The main goal of this article has been to understand egalitarianism. But in these final two sections I will offer some thoughts about its truth.

Suppose that you think that in Myerson's example, L_E is better than L_F . This, I have claimed, is the core egalitarian idea. Then unless you think that the importance of equality is limited to tie-breaking, you should think that it is worth paying some small cost to shift from L_F to L_E . In particular, you should think that for some small $\epsilon > 0$, M , defined below, is better than L_F . This seems fairly obvious, but it would also follow formally if we imposed continuity on theory E.

M	heads	tails	L_F	heads	tails
Alice	$1 - \epsilon$	$-\epsilon$	Alice	1	0
Bob	$1 - \epsilon$	$-\epsilon$	Bob	0	1

But I submit that provided the numbers are properly understood, it is not credible to claim that M is better than L_F . Because L_F strongly dominates M in terms of well-being, in the sense that L_F is better for every individual involved than M , I will call this the *strong levelling-down objection*.

Now because this objection formally relies on continuity, denying continuity might seem to provide egalitarianism with a get-out clause. In response, I suspect that the example would be formally reinstated if we were to move to a richer mathematical model, such as one which admits infinitesimal utilities or probabilities. But I will not pursue this as the move is dialectically disastrous for egalitarians. It tries to avoid

the claim that equality is valueless by starting with the assumption that its value is at most infinitesimal. But no actual egalitarian will be content with that. So from now on I will take continuity for granted.

To state the objection in more detail, I have been silent on the substantive question of when one life is better than another. So there is room within my framework for almost any answer to that question. For example, we could take into account the quality of experience someone's life contains, the extent to which she succeeds in her aims, and her attainment of various objective goods. The numbers would then reflect that, so that a life corresponding to '1' is on balance better in terms of experience, success, and attainment than one corresponding to '1 - ϵ ' and so on.

But once we take into account everything which makes someone's life go well, shifting from L_F to M seems pointless. It makes each person face a lottery $[\frac{1}{2}, 1 - \epsilon; \frac{1}{2}, -\epsilon]$ rather than the strongly dominating $[\frac{1}{2}, 1; \frac{1}{2}, 0]$, but to what purpose? It is double counting to say that the inequality in $[1, 0]$ makes the '1' a worse life than the lives in $[1, 1]$, or the '0' a worse life than the lives in $[0, 0]$. If the inequality in $[1, 0]$ has some kind of bad effect, that has already been taken into account, and the lives with that bad effect included are judged exactly as good as the corresponding lives in $[1, 1]$ and $[0, 0]$.

Now I stress I am not tacitly using an idea from Broome 1991. Broome wished to show that someone who thought that inequality was bad could nevertheless accept anonymous Pareto. To illustrate, suppose we have chosen from among the usual candidates our favoured theory of when one person is better off than another. Suppose also that when the numbers in Myerson's example reflect that theory, we judge that L_E is better than L_F . Broome pointed to a way of modifying the theory of when one person is better off than another, so that the modified theory takes into account not just the goods each person has which are recognized by the original theory (e.g. experience, success, and attainment), but also how well that person is doing in terms of those goods in comparison with others. By using that modification along with a suitable mathematical transform, Broome claimed to have reconciled the original judgement that L_E is better than L_F with anonymous Pareto.

Personally, I doubt this reconciliation is successful. But I am not going to argue that here. My claim is rather that the original judgement that L_E is better than L_F was never plausible in the first place. So if the upside for egalitarians of shifting to a risk-involving framework is that it

solves a number of problems in saying what egalitarianism amounts to, the downside is that it shows that they are vulnerable to what I think is a particularly severe analogue of the levelling-down objection.

Before looking at how egalitarians might respond, let us note two ways in which the strong levelling-down objection is more powerful than Parfit's original levelling-down objection. First, the original objection was that egalitarianism implies that levelling down is an improvement in one respect. But as noted in section 9.4, some have taken it to be an adequate reply that egalitarianism does not imply that levelling down is an overall improvement. This is true when egalitarianism is restricted to histories, for there are moderate versions of egalitarianism. But the present objection is focused around a case in which making everyone worse off to guarantee equality is seen as an overall improvement according to egalitarianism.

Second, the present objection seems to provide a tighter link between egalitarianism's key components. The original objection claims that egalitarianism is subject to a levelling-down objection because of its concern with patterns. But despite that, it offers no concrete claims about the effects of being concerned with patterns on the betterness relation. In particular, it fails to show that being concerned with patterns results in any implausible concrete claims about betterness. In addition, it is often thought that a concern with patterns will lead to the betterness relation violating some sort of separability condition. But the original objection does nothing to connect the levelling-down objection with any such separability violation.

By contrast, the present objection shows how the egalitarian concern with patterns results in a concrete principle about betterness, the equality principle, which then generates an implausible claim about betterness. In addition, the present objection is intimately connected with the egalitarian rejection of weak separability across individual lotteries. It is because they think that shuffling an individual's outcomes between equally likely states of nature can improve patterns without making anyone worse off that egalitarians reject weak separability; but it is that same attitude to shuffling which leads to the strong levelling-down objection.

21. The truth in egalitarianism

How might egalitarians respond to the strong levelling-down objection? One difficulty is that it is not easy to know why many

egalitarians think egalitarianism is correct. Much of the most influential egalitarian literature takes it for granted that egalitarianism is correct in order to focus on more specific questions. In what follows I will examine what I take to be five major ideas which are in some sense egalitarian and which seem to provide support for the idea that equality matters. But I will argue that none of these provides any support for the equality principle. In particular, none rescues it from the strong levelling-down objection. The main contrast in this article has been between egalitarianism and utilitarianism, so that will continue to be my focus.

21.1 Humanitarianism

One natural motivation for egalitarianism is broadly humanitarian. One considers various situations of inequality, finds them unacceptable, and concludes that something like egalitarianism is the right response, usually tempered with enough of a concern with well-being to guarantee that at least when it comes to histories, it is never on-balance better to level down. The debate then focuses on whether egalitarianism is exactly the right response, or whether instead the intuitions are better captured by something like the priority view or a concern with sufficiency. Leading discussions along these lines often do not even mention utilitarianism, and when it is mentioned, it is quickly dismissed.

It is unfortunate that when utilitarianism is characterized at all in these debates, it is usually characterized poorly.¹⁸ So let us look at this style of argument for egalitarianism through the lens of the best argument for utilitarianism, Harsanyi's theorem. As I have characterized it, the dispute between utilitarianism and egalitarianism is over anonymous Pareto versus the equality principle, for they agree about everything else. But while anonymous Pareto entails 'better for everyone is always better', the equality principle implies 'better for everyone is sometimes worse'. But the idea that there is a direct, broadly humanitarian rationale for the latter and against the former is not credible.

But could there be an indirect rationale for the equality principle? To rest one's hopes on this as a defence of egalitarianism already

¹⁸ Just to give one example, one otherwise helpful recent survey tells us that risk is a form of luck, and then confidently informs us that 'no one has argued that utilitarianism is grounded in reflections on luck'. In fact, Harsanyi's theorem, referred to by economists as his utilitarian theorem, is more or less entirely about risk. I mean this less as a criticism of the survey and more as a comment on what it gleans from the large body of egalitarian literature surveyed.

seems a little desperate, for it seems to abandon the idea that equality is intrinsically valuable. On the other hand, an indirect rationale may seem to make the equality principle more palatable as the ‘better for everyone is sometimes worse’ implication may seem like a cost we have to swallow to accommodate some more important principle.

What could that principle be? Perhaps it stems from the kind of case this subsection started with. Let a be a good life, a^+ be a slightly better life, and z a terrible life (severe pain, slavery etc.). Then the kind of allegedly implausible claim to which egalitarianism is supposed to be a reaction is

(C) For a sufficiently large population, $[a^+, \dots, a^+, z]$ is better than $[a, \dots, a, a]$

So perhaps the principle in question is some generalization of the negation of (C). However, the distance from the denial of (C) to the equality principle should not be underestimated. First, egalitarianism is consistent with the acceptance of (C); this follows from Corollary 1. Second, in my view utilitarianism is consistent with the rejection of (C). I lack the space to discuss this, but it follows from my view that the continuity axiom of expected utility theory is not part of the core of utilitarianism.

There is also something odd about trying to defend the equality principle by starting with the negation of (C). Once we see that the equality principle brings the inhumanitarian ‘better for everyone is sometimes worse’ implication in its wake, it seems more natural to explore ways of generalizing the negation of (C) which avoid it. Thus if one is moved to reject (C) for some kind of broadly humanitarian reason, it seems more natural to see this rejection as expressing some sort of antiaggregative idea rather than a concern with equality. I discuss this in McCarthy forthcoming b, but I will not say any more about it here. The denial of (C) is an important idea, and deserves its own day in court.

21.2 *Fairness*

Some writers claim that there is a different rationale for egalitarianism: equality matters because fairness matters (see, for example, Broome 1991, Temkin 1993, Otsuka 2004). To discuss this I need to abandon the way I have been talking about fairness and revert to my favoured terminology, according to which Myerson’s example is about equality of outcome and Diamond’s example is about equality of prospects. This terminology is helpfully neutral because it is silent on how

fairness relates to either of those forms of equality. The view we are considering then says that equality of outcome matters because fairness matters.

Let us look at this more closely. Nothing in the ordinary use of 'fair' forces us to say that the presence of equality guarantees fairness while its absence guarantees unfairness. First, suppose A and B have to choose between an outcome [10, 10] and a lottery which will almost certainly result in [100, 100] but involves a tiny chance of [0, 0]. They choose the lottery. But they are unlucky and end up with [0, 0], and bemoan their fate by complaining that life is unfair. Second, consider the view that because L_F distributes the chances perfectly equally, there is nothing unfair about whichever outcome results, despite the inequality it contains. Neither of these usages of 'unfair' is a misuse of the term, so evidently our ordinary usage of the term is quite flexible. As Rawls (1971) might have put it, there are different conceptions of fairness. This may be what leads Temkin (2003), for example, to claim that the fundamental egalitarian concern is with what he calls 'comparative fairness'.

I take comparative fairness to be a particular conception of fairness, but what exactly does it amount to? There is said to be comparative unfairness when there is inequality of outcome. Temkin thinks that unfairness makes situations worse in one respect, and if the arguments of this article are along the right lines, the concrete manifestation of this idea is just the equality principle. But it is now clear that there is not enough distance between the two to say that equality of outcome matters because fairness matters.

I suspect that those who claim that equality of outcome matters because fairness matters are making a subtle mistake about the concept of fairness. For very roughly, I suspect that 'unfair' is typically used to signal a departure from some background value or norm, with focus on those who are worse off as a result. For example, utilitarians could claim that people are treated unfairly when their well-being is given less weight than everyone else's; libertarians could claim that people are treated unfairly when some of the product of their labour is redirected towards the worse off; and so on. But in this usage, applications of 'unfair' are clearly parasitic upon the acceptance of the background value or norm.¹⁹ One might call this a *free-floating* theory of fairness. So it is plainly no argument for the value or

¹⁹ This is why I do not say that Diamond's example is about fairness. More or less any putative distributive value carries with it its own conception of unfairness.

norm to claim that a particular type of situation contains unfairness. In particular, the claim that equality of outcome matters because fairness matters provides us with no independent reason to accept the equality principle, and thus provides no independent support for the disturbing implication ‘better for everyone is sometimes worse’.

21.3 The separateness of persons

Rawls (1971) famously argued that utilitarianism improperly ignores the separateness of persons. This may seem to suggest an argument for egalitarianism. For example, Parfit writes

The argument for equality is often claimed to be grounded on the separateness of persons. One such claim might be: ‘Since it is a deep truth that we lead different lives, it is an ultimate moral aim that ... the lives of each should go equally well.’ (Parfit 1984, p. 330)

This line of argument for egalitarianism is doubly mistaken. Rawls’s criticism of utilitarianism took it for granted that utilitarianism is based on the impartial spectator argument. At the heart of this argument is the claim that one history is better than another if and only if a rational and self-interested spectator would strictly prefer fully experiencing every life in the first history to fully experiencing every life in the second. Rawls (1971, p. 27) claimed that this approach improperly ignores the boundaries between lives. Parfit (1984, pp. 330 ff.) replied that utilitarians should instead be understood as claiming that the boundaries between lives have no moral significance, because they accept what he called a reductionist view about personal identity. However, it is not so clear how far this response goes towards supporting utilitarianism. The boundaries between lives are still recognizable on a reductionist view, and as Parfit himself notes, it is hard to see why utilitarianism emerges from his appeal to personal identity as plausible rather than as relatively less implausible.

A better response to Rawls is simply to point out that he overlooked an argument for utilitarianism which is invulnerable to his objection. But more or less following Harsanyi, what I called the utilitarian premisses, listed in section 8.2, are clearly subject to nothing like Rawls’s separateness of persons objection. Rawls therefore failed to identify a defect in utilitarianism to which egalitarianism can be seen as offering a remedy. If egalitarianism is going to appeal to the separateness of persons, it is going to have to do so directly. But it is vastly implausible that anything like an appeal to the fact that we lead

different lives should support the claim that in Myerson's example, L_E is better than L_F .

21.4 *Equality, surely!*

It seems obvious that people are equal in some fundamental respect. But then any plausible account of the ethics of distribution should reflect that fundamental fact. But how better to do that than by adopting the idea that distributive equality is valuable, and hence the equality principle?

In response, once we restrict our attention to questions about distribution, by far the most obvious respect in which people are equal is captured by the risk-free version of anonymous Pareto: A well off and B badly off is exactly as good as A badly off and B well off, no matter what A's and B's individual characteristics are, such as ethnicity, gender, sexuality, and the like. The fact that this is a commonplace in discussions of the ethics of distribution should not disguise the fact that it expresses a profound sense in which people are equal, and historically is arguably the kind of equality which many egalitarian movements have fought for (cf. Anderson 1999). But obviously, this is no motivation for the equality principle, for risk-free anonymous Pareto is consistent with every distributive view we have been looking at, including utilitarianism.

21.5 *Egalitarian goods*

Thanks largely to the egalitarian literature, we now have a much wider range of candidate answers to the question of what is good for people. Some of these answers may seem to support egalitarianism. To illustrate, Sen (1980) and Rawls (1982) criticize utilitarianism because it assumes what Sen calls *welfarism*, the claim that what is good for people is pleasure or desire satisfaction. In place of welfarism they and many other egalitarians recommend understanding what is good for people in part in terms of various goods to do with freedom and nonsubordination. Moreover, some of these goods are relational in a way which, without trying to be too precise about it, makes them 'equality-involving'. For example, if you are not subordinate to others, others do not have corresponding powers over you. Thus one might be tempted to conclude that recognizing the importance of these goods provides support for egalitarianism over utilitarianism, and thus for the equality principle. For utilitarianism cannot acknowledge the goods, and the equality-involving nature of some of them might seem to provide direct support for the equality principle.

But this impression should not last for long. Sen and Rawls are really discussing the question of how should we make interpersonal (and intrapersonal) comparisons; or in my terminology, the question of what the content of the extended individual betterness relation is. This is a compulsory question for all distributive views, and it is hard to see what would override the *prima facie* independence of the two questions: What should we care about distributing, and how should we distribute it? To illustrate, regardless of any historical association with welfarism, utilitarianism need not be tied to it. Utilitarianism is better characterized as offering an answer to the question: Given an account of the content of the extended individual betterness relation, what is the content of the betterness relation? Utilitarianism can help itself to any view about the content of the former, including those views developed by egalitarians. In particular, if freedoms are indeed important goods, this will be factored into the content of the extended individual betterness relation. The numbers in Myerson's example then encode people's holdings of all goods, the freedoms included. But then, to repeat, making everyone worse off in terms of these goods, equality-involving or not, is pointless.²⁰

References

- Anderson, Elizabeth 1999: 'What is the Point of Equality?' *Ethics*, 109, pp. 287–65.
- Arneson, Richard 2001: 'Luck and Equality'. *Proceedings of the Aristotelian Society, Supplementary Volume*, 75, pp. 73–90.
- Atkinson, Anthony B. 1970: 'On the Measurement of Inequality'. *Journal of Economic Theory*, 2, pp. 244–63.
- Ben-Porath, Elchanan, Itzhak Gilboa, and David Schmeidler 1997: 'On the Measurement of Inequality under Uncertainty'. *Journal of Economic Theory*, 75, pp. 194–204.

²⁰ An early version of this article was given at the LSE in 2007, and a more recent version at Edinburgh in 2011. I thank members of the audiences for comments, and I especially thank Richard Bradley, John Broome, Branden Fitelson, Marc Fleurbaey, Natalie Gold, Peter Milne, and Alex Voorhoeve. I have benefited greatly from extensive comments on multiple drafts by three referees and the editor for *Mind*. I am hugely grateful to Teru Thomas for joint work strengthening a result of the early version. This joint work is in part a formal development of the present article, and though I have tried to minimize overlap, preparation of the final version of this article has no doubt greatly benefitted from it. Support was partially provided by the AHRC Research Leave Scheme and by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (HKU 750012H).

- Bosmans, Kristof 2007: 'Extreme Inequality Aversion without Separability'. *Economic Theory*, 32, pp. 589–94.
- Broome, John 1989: 'What's the Good of Equality?' In Hey 1989, pp. 236–62.
- 1990–91: 'Fairness'. *Proceedings of the Aristotelian Society*, 91, pp. 87–102.
- 1991: *Weighing Goods*. Oxford: Blackwell.
- 2004: *Weighing Lives*. New York: Oxford University Press.
- forthcoming: 'Equality versus Priority: A Useful Distinction'. *Economics and Philosophy*.
- Clayton, Matthew and Andrew Williams (eds) 2000: *The Ideal of Equality*. Basingstoke: Macmillan.
- Cohen, G. A. 1989: 'On the Currency of Egalitarian Justice'. *Ethics*, 99, pp. 906–44.
- Diamond, Peter 1967: 'Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility: Comment'. *Journal of Political Economy*, 75, pp. 765–6.
- Dworkin, Ronald 2000: *Sovereign Virtue*, Cambridge MA: Harvard University Press.
- Egonsson, Dan, Jonas Josefsson, Bjorn Petersson, Toni Rønnow-Rasmussen, and Ingmar Persson (eds) 2001: *Exploring Practical Philosophy: From Action to Values*. Aldershot: Ashgate.
- Elster, Jon and John E. Roemer (eds) 1991: *Interpersonal Comparisons of Well-Being*. New York: Cambridge University Press.
- Fleurbaey, Marc 2010: 'Assessing Risky Social Situations'. *Journal of Political Economy*, 118, pp. 649–80.
- forthcoming: 'Equality versus Priority: How Relevant is the Distinction?' *Economics and Philosophy*.
- Hájek, Alan and Christopher Hitchcock (eds) forthcoming: *The Oxford Handbook of Philosophy and Probability*. Oxford University Press.
- Hammond, Peter 1976: 'Equity, Arrow's conditions, and Rawls' Difference Principle'. *Econometrica*, 44, pp. 793–804.
- 1991: 'Interpersonal Comparisons of Utility: Why and How They Are and Should Be Made'. In Elster and Roemer 1991, pp. 200–54.
- Harel, Alon, Zvi Safra, and Uzi Segal 2005: 'Ex-post Egalitarianism and Legal Justice'. *The Journal of Law, Economics, and Organization*, 21, pp. 57–75.
- Harsanyi, John 1955: 'Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility'. *Journal of Political Economy*, 63, pp. 309–21.

- Hey, John D. (ed.) 1989: *Current Issues in Microeconomics*. Basingstoke: Macmillan.
- Hurley, Susan 2001: 'Luck and Equality'. *Proceedings of the Aristotelian Society, Supplementary Volume*, 75, pp. 51–72.
- Jensen, Karsten Klint 2003: 'What is the Difference between (Moderate) Egalitarianism and Prioritarianism?' *Economics and Philosophy*, 19, pp. 89–109.
- Krantz, David H., R. Duncan Luce, Patrick Suppes, and Amos Tversky 1971: *Foundations of Measurement*, Vol. 1. New York: Academic Press.
- Kreps, David 1988: *Notes on the Theory of Choice*. Colorado: Westview Press.
- McCarthy, David 2007: 'Measuring Life's Goodness'. *Philosophical Books*, 48, pp. 303–19.
- 2008: 'Utilitarianism and Prioritarianism II'. *Economics and Philosophy*, 24, pp. 1–33.
- 2013: 'Risk-free Approaches to the Priority View'. *Erkenntnis*, 78, pp. 421–49. doi:10.1007/s10670-012-9377-4.
- forthcoming a: 'Probability in Ethics'. In Hájek and C. Hitchcock forthcoming.
- forthcoming b: *The Structure of Good*. Oxford University Press.
- forthcoming c: 'The Priority View', forthcoming in *Economics and Philosophy*.
- McCarthy, David, Kalle Mikkola, and Teruji Thomas MS: 'Utilitarianism with and without Expected Utility'.
- McCarthy, David and Teruji Thomas MS: 'Egalitarianism with Risk'.
- McKerlie, Dennis 1989: 'Equality and Time'. *Ethics*, 99, pp. 475–91.
- McMurrin, Stirling M. (ed.) 1980: *The Tanner Lecture on Human Values*, Vol. I. Cambridge: Cambridge University Press.
- Myerson, Roger B. 1981: 'Utilitarianism, Egalitarianism, and the Timing Effect in Social Choice Problems'. *Econometrica*, 49, pp. 883–97.
- Narens, Louis 2007: *Introduction to the Theories of Measurement and Meaningfulness and the Use of Symmetry in Science*. Mahwah, NJ: Lawrence Erlbaum.
- Ok, Efe 2007: *Real Analysis with Economic Applications*. Princeton: Princeton University Press.
- Otsuka, Michael 2004: 'Equality, Ambition, and Insurance'. *Aristotelian Society, Supplementary Volume*, 78, pp. 151–66.
- Otsuka, Michael and Alex Voorhoeve 2009: 'Why it Matters that Some are Worse than Others: An Argument against the Priority View'. *Philosophy and Public Affairs*, 37, pp. 171–99.

- Parfit, Derek 1984: *Reasons and Persons*. New York: Oxford University Press.
- 2000: 'Equality or Priority?'. Clayton and Williams 2000, pp. 81–125.
- Rabinowicz, Wlodek 2001: 'Prioritarianism and Uncertainty: On the Interpersonal Addition Theorem and the Priority View'. In Egonsson et al. 2001.
- 2002: 'Prioritarianism for Prospects'. *Utilitas*, 14, pp. 2–21.
- Rawls, John 1971: *A Theory of Justice*. Cambridge MA: Harvard University Press.
- 1982: 'Social Unity and Primary Goods'. Sen and Williams 1982, pp. 159–85.
- Resnik, Michael D. 1987: *Choices: An Introduction to Decision Theory*. Minneapolis: University of Minnesota Press.
- Roberts, Fred S. 2009: *Measurement Theory*. Cambridge: Cambridge University Press.
- Segal, Uzi 2006: 'Fair Bias'. *Economics and Philosophy*, 22, pp. 213–29.
- Sen, Amartya 1973: *On Economic Inequality*. Oxford: Clarendon Press.
- 1980: 'Equality of What?' In McMurrin 1980, pp. 195–220.
- Sen, Amartya and Bernard Williams 1982: *Utilitarianism and Beyond*. Cambridge: Cambridge University Press.
- Temkin, Larry S. 1993: *Inequality*. New York: Oxford University Press.
- 2003: 'Exploring the Roots of Egalitarian Concerns'. *Theoria*, 69, pp. 125–51.
- Tungodden, Bertil 2000: 'Egalitarianism: Is Leximin the Only Option?' *Economics and Philosophy*, 16, pp. 229–45.
- 2003: 'The Value of Equality'. *Economics and Philosophy*, 19, pp. 1–44.
- Tungodden, Bertil and Peter Vallentyne 2005: 'On the Possibility of Paretian Egalitarianism'. *Journal of Philosophy*, 102, pp. 126–54.
- Vallentyne, Peter 2000: 'Equality, Efficiency, and Priority of the Worse-off'. *Economics and Philosophy*, 16, pp. 1–19.

