Causal bias in measures of inequality of opportunity

Lennart Ackermans

September 15, 2021

1 Introduction

Recent decades have seen a surge in interest to quantitatively measure inequality of opportunity.[1] Several studies have come out estimating the amount of inequality of opportunity in a country (Bourguignon, Ferreira, and Menéndez 2007; Checchi and Peragine 2010; Pistolesi 2008; Almås et al. 2011; Björklund, Jäntti, and Roemer 2012; Davillas and Jones 2020) and comparing different countries with respect to inequality of opportunity (Lefranc, Pistolesi, and Trannoy 2008; Ferreira and Gignoux 2011; Checchi, Peragine, and Serlenga 2016; Hufe et al. 2017; Brunori, Palmisano, and Peragine 2019).

These studies use one of two methodological approaches, the *ex-ante* and *ex-post* approach. The *ex-ante* approach focuses on inequality between groups of people who share the same circumstances. Such an approach could be understood as measuring the undeserved effects of circumstances on unequal outcomes, or it could measure the value of the opportunity set that all people within some group are assumed to face. The *ex-post* approach focuses on inequality between individuals who are the same in matters of responsibility (e.g. they made similar choices) but have different circumstances. This approach could be understood as measuring to what extent inequalities are the result of individual choices and to what extent they are not. Which factors classify as circumstance and responsibility is up to the authors of each study; it will not be of concern in this paper.

I identify serious problems with *ex-ante* measures of inequality of opportunity, which is the most popular method in empirical studies. For *ex-ante* methods to produce reliable results that measure normatively appropriate types of inequality of opportunity, they should identify causal parameters rather than mere statistical associations, but the methods are not designed to do that. In order to show this for two versions of the *ex-ante* approach (parametric and non-parametric), I identify the type of inequality of opportunity that the respective methods aim to measure and show how they fail to do so in the presence of unmeasured confounding factors. (I attempt to stay close to the stated aims of the studies that use these methods, but at same time interpret their definitions

[1]For an overview of the literature, see Ramos and Van de gaer (2016), Roemer and Trannoy (2015), and Ferreira and Peragine (2016).

of inequality of opportunity such that they are clear and normatively viable. Studies need to both measure something important as well as do so successfully.)

While the paper is mainly concerned with the *ex-ante* approach, a subset of *ex-post* methods that use what is called Roemer's Identification Assumption face similar problems, on which I briefly reflect in the conclusion.

I start by showing how these problems affect the *parametric ex-ante* approach (section 3). Parametric methods rely on estimating the contribution of circumstances to outcomes, which they use to create a counterfactual distribution supposed to arise if everyone were equal in terms of circumstances. Inequality of opportunity is thought to be the total amount of inequality minus the counterfactual inequality if circumstances were equal.

However, these methods estimate the contribution of circumstances to outcomes using standard regression analysis on observational data. I use results from the causal modeling literature (especially Pearl 2009) to argue that coefficients estimated by the parametric *ex-ante* approach cannot be given a causal interpretation, and that it is nearly impossible to correct the method using typical back-door adjustment on confounding factors. I further show that, as a result, the estimated amount of inequality of opportunity is unlikely to reflect the true amount of inequality of opportunity.

The non-parametric *ex-ante* approach faces a similar problem (section 5). This approach intends to measure the value of opportunity sets. Opportunity sets, however, are not directly measurable, so non-parametric *ex-ante* methods attempt to estimate the value of opportunity sets from the outcome distributions conditional on circumstances. Individuals who share the same circumstances are thought of as a *type*. A type's distribution of outcomes is assumed to reflect the value of that type's opportunity set.

I show that the non-parametric *ex-ante* approach produces unreliable results under essentially the same conditions as the parametric *ex-ante* approach, even if the former is thought of as measuring a distinct quantity from the latter (inequality of opportunity sets rather than inequality caused by circumstances). When circumstances are confounded by unmeasured variables, responsibility variables will not be equally distributed across types. This leads to a biased estimation of the value of opportunity sets based on the type's outcome distribution.

Other sources of bias have been identified. Many authors recognize that inequality of opportunity measures should be interpreted as lower bounds, since the exclusion of unobserved circumstances would lead to significant downward bias (Ferreira and Gignoux 2011; Lara Ibarra and Martinez-Cruz 2015; Juárez 2015; Balcázar 2015). A recent paper argues that there may also be upward bias as a result of sampling variance (Brunori, Palmisano, and Peragine 2019). While the results in this paper are similar to those showing downward bias due to omitted variables, I give a new causal analysis of the problem, centered around the notion of confounding bias. Moreover, I show that this type of bias can also be upwards. Since the extent of this bias cannot be estimated, existing inequality of opportunity measures might be too biased to be useful.

Section 2 introduces the *ex-ante* and *ex-post* approaches to measuring inequality of opportunity. Section 3 shows that the parametric *ex-ante* approach is causally biased, using insights from causal modeling. Section 4 argues that the problem is severe, and responds to the objection that the problem may be less severe under appropriate classifications of circumstance and responsibility factors. Section 5 shows that the non-parametric *ex-ante* approach, understood as measuring opportunity sets, faces similar problems as the parametric approach. Section 6 concludes and gives some thoughts about future research.

2 Measurement approaches and normative foundations

While the measurement literature stands largely on its own, it is heavily inspired by works of philosophers in the tradition of luck egalitarianism, such as Cohen (1989), Dworkin (2002), and Arneson (1989), among others. For overviews of both the measurement approaches and normative foundations, see Ramos and Van de gaer (2016), Roemer and Trannoy (2015), and Ferreira and Peragine (2016). In this section I summarize what I think are the three normative principles underlying various measurement approaches of inequality of opportunity, and outline the important distinctions between these approaches.

As the principle is typically expressed in the economic literature, equality of opportunity obtains when differential outcomes are the result of factors for which individuals are responsible, but not the result of circumstances, for which individuals are not responsible. This definition comprises two elements: the first, which we may call *rewards to effort*, is that only inequalities that are the result of choices for which individuals are responsible are deemed acceptable; the second, which we may call *equality of luck*, is that inequalities that are caused by circumstances – typically defined as factors outside of individuals' control – are deemed unacceptable. A third definition, introduced by Van de gaer (1993), is that each individual should face the same set of opportunities. Opportunity sets consist of combination of effort choices and associated outcomes that individuals are able to choose from. Call this *equality of opportunity sets*.

To give full content to these three principles one needs to categorize factors as responsibility or circumstances variables, a question that is not of concern in this paper. Another way in which these approaches differ is in the outcome variable that is considered. Many studies are concerned with income inequality (of opportunity), but other outcomes such as health outcomes have also been considered.

It should be noted that the word 'effort' is used as a shorthand for the combined matters of responsibility. Effort is typically a scalar variable that is assumed to reflect all factors that individuals are responsible for. Effort in the usual sense of the term (exertion) may but does not need to be one of these factors.

The economic literature identifies two broad categories of measurement approaches, the *ex-ante* and *ex-post* approach (Fleurbaey and Peragine 2013). The *ex-ante* approach focuses on inequality between *types*, which are groups of individuals who share the same circumstances. The *ex-post* approach focuses on inequality between *tranches*, which are groups of individuals who are the same in matters of responsibility. While these two approaches are frequently formulated in such a way that they could be seen as normative principles of equality of opportunity, they should rather be understood as different methodological approaches to measuring inequality of opportunity. The approaches are distinguished both by different (but not necessarily conflicting) normative intuitions as well as different empirical assumptions.

Ex-ante inequality of opportunity can be defined as follows.

> **Ex-ante inequality of opportunity:** Let $\mu(t)$ be a measure of the advantage of type t. Inequality of opportunity decreases if inequality between types, $\mu(t_1) - \mu(t_2)$, decreases.

Studies taking the *ex-ante* approach differ with respect to the measure $\mu(t)$ they adopt. Most commonly, $\mu(t)$ is simply taken to be the average outcome of individuals within type t. The measure $\mu(t)$ is most often thought of as measuring the value of a type's opportunity set, in which case the normative justification of the measurement approach based on equality of opportunity sets (see section 5). Alternatively, $\mu(t)$ can be thought of as the causal effect of the type's circumstances on outcomes, in which case the normative justification is based on equality of luck (see section 3).

Ex-post inequality of opportunity can be defined as follows.

> **Ex-post inequality of opportunity:** Let $f(e, t)$ be the outcome of an individual with effort level e and type t. Inequality of opportunity decreases if outcome inequality between individuals with the same effort level, $f(e, t_1) - f(e, t_2)$, decreases.

Studies taking the *ex-post* approach differ with respect to the way in which inequality within different tranches is aggregated (or which tranches are considered in the first place). Importantly, they also differ in the way that effort (e) is measured.

Ex-post approaches face the problem that effort is typically hard to measure. A 'solution' to this problem, adopted by many studies, is to measure effort indirectly via Roemer's Identification Assumption (RIA), which identifies an individual's effort with their position in their type's outcome distribution. According to Roemer (2002) this assumption ensures that we take into account the effect of circumstances on actual (observed) effort, since we hold people responsible only for their relative effort level within their type. When using RIA, the *ex-post* approach becomes similar to the *ex-ante* approach, and in some cases identical (see also Ramos and Van de gaer 2016). For example, one could use RIA and choose to measure inequality only in the median effort tranche. This is identical to an *ex-ante* approach in which $\mu(t)$ is chosen to be the median income of type t. The rest of the paper focuses on the *ex-ante* approach, but in the conclusion I briefly reflect on *ex-post* approaches using RIA, to which the arguments in the paper apply as well.

3 Measuring the causal effect of circumstances on outcomes

Bourguignon, Ferreira, and Menéndez (2007) introduce an approach to measuring inequality of opportunity based on an estimation of the (causal) contribution of observed circumstances to wages, and apply this approach to measure inequality of opportunity in Brazil.[2] Since this approach is centered around estimating parameters that reflect

[2]Due to a programming error the estimates in Bourguignon, Ferreira, and Menéndez (2007) are faulty. See the corrections in Bourguignon, Ferreira, and Menéndez (2013).

the (causal) contribution of circumstances to outcomes, it is called a *parametric ex-ante* approach. Most *ex-ante* approaches are non-parametric, but Ferreira and Gignoux (2011) show that the parametric estimate of inequality of opportunity discussed in this section can be interpreted as the same quantity (measured differently) as the non-parametric estimate that will be discussed below in section 5.

The normative foundation of the parametric ex-ante approach as used by Bourguignon, Ferreira, and Menéndez (2007) is *equality of luck*, the principle that circumstances should not cause unequal outcomes. In this spirit, Bourguignon et al. set out to measure not just the direct effect of circumstances on outcomes, but also their indirect effect on outcomes via their effect on effort. Equality of opportunity is said to occur if the total effect (the combined direct and indirect effect) is the same for each type.

The estimates are made by OLS on

$$\ln(w_i) = \psi C_i + \varepsilon_i. \tag{1}$$

Here the regressand w_i is the wage of the i'th individual, the regressor C_i is a vector of observed circumstance values (plus a constant), ψ is a vector of coefficients and ε_i an error term with mean 0. Note that equation (1) does not contain effort variables. The part of effort that is not caused by circumstances is assumed to be represented in the error term ε_i. (See Bourguignon, Ferreira, and Menéndez 2007 for a derivation of (1) from equations containing effort variables.)

Based on the estimated coefficients $\hat{\psi}$ and error terms $\hat{\varepsilon}_i$, Bourguignon et al. calculate a counterfactual earnings distribution X^C, based on the model $\ln \tilde{w}_i = \bar{C}\hat{\psi}_i + \hat{\varepsilon}_i$. Here $\bar{C}$ denotes the population mean of the circumstance variables. The distribution X^C is interpreted as the distribution that would arise if everyone had the same circumstances, and is therefore a situation in which there is equality of opportunity (understood as equality of luck). The counterfactual earnings distribution is compared to the actual earnings distribution X to obtain a measure of inequality of opportunity:

$$\Theta_I := I(X) - I(X^C) \tag{2}$$

Here I denotes an inequality index. Borguignon et al. use the Theil index for their own calculations.

As I show in example 3.1 below, the coefficients ψ need to be given a causal interpretation if they are used to measure inequality of opportunity. However, regression coefficients only represent causal contributions under strict conditions. In general, a causal interpretation requires that the measured variables (C) are not confounded by unobserved variables (see e.g. Pearl 2009). Confounding would occur if there is a common cause of a circumstance and wages that is not adjusted for by other regressors.

Bourguignon, Ferreira, and Menéndez (2007) are not unaware of such problems. They note that the estimated coefficients are biased if the circumstance variables C are econometrically exogenous. Econometric exogeneity of C means that ε and C are uncorrelated and that ε has mean 0, such that $\mathbf{E}[\varepsilon C] = 0$. They propose to explore the likely magnitude of this bias using a Monte-Carlo method that considers a wide range of estimates of the bias given random draws of values for the correlation between variables and error terms (assuming uniform distributions on $[-1, 1]$). This procedure seems to lead to reasonably precise lower bounds for the true inequality of opportunity (Bourguignon, Ferreira, and Menéndez 2013). However, the procedure only explores bias due to a lack of econometric exogeneity, whereas bias due to confounding is a larger problem. The example below, in which econometric exogeneity and confounding co-occur, illustrates the problem.

Similarly, Ferreira and Gignoux (2011) proof that (2) underestimates inequality of opportunity if not all circumstances are measured, assuming that circumstances may have an effect on effort but not the other way around. However, this is so only in an econometric sense: econometric unbiasedness, which follows from econometric exogeneity, means that $\mathbf{E}[\hat{\psi}] = \psi$, where $\hat{\psi}$ is the OLS estimator of ψ. Confounding, however, could imply that ψ itself does not represent the causal contribution of circumstances. If circumstances and effort variables have commons causes, then (2) may also overestimate the true inequality of opportunity, as the example below demonstrates.

3.1 Example: common causes

The following example serves to illustrate both that econometric exogeneity is insufficient to measure causal parameters, as well as that a normatively acceptable measure of inequality of opportunity needs to measure causal, rather than merely econometric, parameters. In other words, this is an example of confounding bias.

We examine a population of individuals who go to separate schools of different levels, spend some time on schooling, and then enter the labor market. We seek to estimate the effect of the school's level – considered a circumstance – on wages, in order to measure income inequality of opportunity.

Suppose that individuals in our population differ with respect to their ambition (A), which we consider an effort variable. For mathematical simplicity, suppose that school level (S) is a continuous variable. Individuals are assigned to a school based on their ambition according to

$$S_i = rA_i, \tag{3}$$

where r is some unknown constant.

After finishing their education, individuals enter the labor market. We suppose that their productivity P is determined linearly from A as

$$P_i = sA_i, \tag{4}$$

for some constant s. Employers set wages according to

$$w_i = \alpha S_i + \beta P_i + u_i. \tag{5}$$

Here α is the causal contribution of school level to wages, and β the causal contribution of productivity to wages. We suppose that u is normally distributed about 0 and uncorrelated with S and P, such that we have $E[Su] = 0$ and $E[Pu] = 0$. The full structural model is depicted as a causal graph in figure 1.

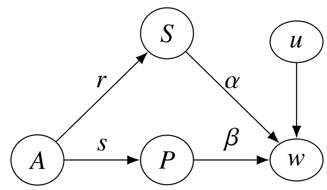

Figure 1: Causal graph of example 3.1.

Suppose the researcher is unaware of the way in which S and P are determined from A, and is able to measure only S. She proposes to estimate the contribution of S to w by regressing

$$w_i = \psi S_i + \varepsilon_i. \tag{6}$$

By consecutive substitution of (4) and (3) into (5), we can write the causal model as

$$w_i = (\alpha + \frac{\beta s}{r})S_i + u_i. \tag{7}$$

Hence, the omitted variables in the regression of (6) will be completely absorbed by ψ rather than ε_i, and we have $\varepsilon_i = u_i$. Hence, $\mathbf{E}[S\varepsilon] = \mathbf{E}[Su] = 0$, so S in (6) is econometrically exogenous. It follows that the expected value of the OLS estimator is $\mathbf{E}[\hat{\psi}] = \alpha + \beta s/r$. While $\hat{\psi}$ is an unbiased estimate of ψ, it is a biased estimate of α, the causal effect of S on w.

It is easy to see that the quantity of interest is indeed α and not $\alpha + \beta s/r$. A situation in which equality of opportunity occurs – given the intended interpretation as equality of luck – is when school assignment does not cause inequalities. This situation would occur if $\alpha = 0$, but then the regression of (6) would lead to an expected estimate of $\mathbf{E}[\hat{\psi}] = \beta s/r$.

The following shows that Θ_I is a biased estimator of the true inequality of opportunity. In the situation that everyone would be assigned the same school level $\bar{S}$, that is, equation (3) would be replaced by $S_i = \bar{S}$, then wages would satisfy

$$w_i = \alpha\bar{S} + \beta P_i + u_i. \tag{8}$$

An estimator of inequality of opportunity should measure the difference between the actual distribution generated by (5) and the counterfactual distribution generated by (8), which is the distribution that would arise if circumstances were equal. The estimator Θ_I, on the other hand, measures the difference between the distribution generated by (5) and

$$w_i = (\alpha + \frac{\beta s}{r})\bar{S} + u_i. \tag{9}$$

It is not clear how this distribution can be interpreted as a counterfactual distribution at all. Since $(\alpha + \beta s/r)\bar{S}$ is a constant, a distribution generated by (9) contains only inequalities that are the result of the random disturbances u_i. The appropriate counterfactual distribution generated by (8), on the other hand, also contains inequalities that are the result of differences in productivity P (which in turn are the result of differences in ambition). Hence, in this example, Θ_I overestimates inequality of opportunity.

3.2 Can we adjust for confounding?

One might wonder whether it is possible to improve upon the expounded methodology in order to adjust for or reduce confounding bias. While such adjustment methods exist, they would achieve little with the limited observational data that most studies of (income) inequality of opportunity have been using so far.

A criterion for sufficiently adjusting for confounding called the back-door criterion is given by Pearl (2009). Application of this criterion requires knowledge of the causal mechanism that is described by a causal model using a directed acyclic graph (DAG). To measure the causal contribution of a variable X to Y one needs a set of adjustment variables Z that satisfy two conditions (the back-door criterion below).

First, we need some definitions. In a DAG, a path of three variables is called a chain if they are connected by arrows that go in one direction, such as the chain $A \rightarrow S \rightarrow w$ in figure 1. It is called a fork if the variables move out from the middle variable, such as the fork $S \leftarrow A \rightarrow P$ in figure 1. It is called a collider if the arrows point to middle variable, such as the collider $P \rightarrow w \leftarrow S$ in figure 1. When adjusting using the back-door criterion, confounding causal paths need to be *blocked* by a set of adjustment variables. Blocking is defined as follows.

> ***d*-separation or blocking (Pearl 2009):** a path p is blocked by a set of variables Z if and only if
>
> (i) p contains a chain $i \rightarrow m \rightarrow j$ or a fork $i \leftarrow m \rightarrow j$ such that the middle node m is in Z; or
>
> (ii) p contains a collider $i \rightarrow m \leftarrow j$ such that the middle node m is not in Z and such that no descendant of m is in Z.

One adjusts for confounding using the back-door criterion by conditioning on the right set Z. In a linear regression model, this procedure is carried out by adding these variables as regressors. In such a case, the set of regressors Z other than the causal variable of interest, X, must satisfy the back-door criterion. Then the partial regression coefficient of X conditional on Z is a reliable estimate of the total causal effect of X on the dependent variable, Y (152).

> **Back-door criterion (Pearl 2009):** Z satisfies the back-door criterion relative to a cause X, effect Y and a DAG G if
>
> (i) no node of Z is a descendant of X; and
>
> (ii) Z blocks every path between X and Y that contains an arrow into X.

Condition (i) ensures that we don't condition on the wrong variables, which lie on the causal path we want to estimate. Condition (ii) requires that variables on the path between a common cause, X and Y are included in Z to adjust for.

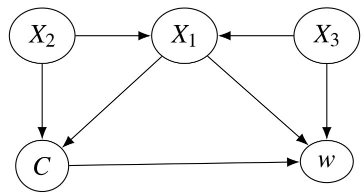

Figure 2: Example in which conditioning on X_1 blocks and unblocks a path at the same time.

Suppose we want to measure the causal effect of P on w in figure 1. Both $Z = \{A\}$ and $Z = \{A, S\}$ satisfy the back-door criterion. $Z = \emptyset$ does not, since the path $P \leftarrow A \rightarrow S \rightarrow w$ would not be blocked.

It is in theory possible to estimate the causal contribution of C to income using back-door adjustment. However, to execute this procedure correctly, one needs extensive knowledge of the underlying causal mechanism and comprehensive data to make the required adjustments. Specifically, one needs to have information about each path between C and w via a common cause. Such paths are likely plentiful (see section 4), and it seems impossible that we could ever unveil the full structure, let alone gather data for each of these paths.

Even if we had data for all relevant variables, it would be easy to mistakenly condition on the wrong variables if the full causal model is not known. This could happen when one conditions on a collider, which could unblock paths that would otherwise be blocked. Sometimes, however, conditioning on a collider is necessary. Consider a case in which a common cause X_1 of C and w is also a collider in a different path between C and w, as in figure 2. In this model, conditioning on X_1 is necessary to block the path $C \leftarrow X_1 \rightarrow w$, and it would also block the other previously unblocked paths; but at the same time, adjustment for X_1 would unblock the path $C \leftarrow X_2 \rightarrow X_1 \leftarrow X_3 \rightarrow w$, in which it is a collider. Hence, one should additionally condition on X_2, X_3, or both. Given the complexity of the real causal mechanism behind wages, it is likely that the selection of suitable adjustment variables is a non-trivial matter for which one needs knowledge of the causal mechanism, as in the example.

A related problem occurs if the data itself is conditional on a collider. This problem could be called sample selection bias; but in our case, it is likely to occur even if samples are representative of the population, as will be illustrated in example 3.3 below.

When we have incomplete knowledge of the causal mechanism and limited data, causal effects are identifiable only under special conditions, such as when the available data allows for an instrumental variable approach. Existing studies of inequality of opportunity all use observational data, and none uses an instrumental variables approach. Possibly, existing data does not allow the causal effect of circumstances on wages to be identified at all.

3.3 Example: 'sample selection' bias

The following example gives a different way in which inequality of opportunity might be overestimated due to the regression coefficients not matching the causal effects.

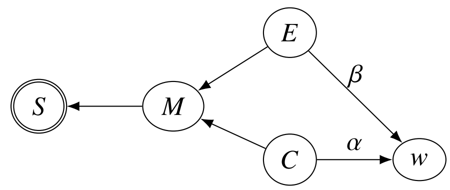

Figure 3: Causal graph of example 3.3. The node M indicates migration. Following Bareinboim, Tian, and Pearl (2014), the sample selection node S is given two circles.

Existing studies of inequality of opportunity use large data sets that might be argued to be representative of the population, as do Bourguignon, Ferreira, and Menéndez (2007). However, a similar sort of problem is that the population itself has been altered – by death and migration – in a way that creates a correlation between circumstances and effort. The structure of this problem is the same as ordinary sample selection bias. See Bareinboim, Tian, and Pearl (2014) and Bareinboim and Pearl (2016) for a discussion of sample selection bias in the context of measuring causal effects.

Suppose we have one circumstance variable C, and one effort variable E, which are initially independently normally distributed. A lower value of C means a worse circumstance, and lower value of E means lower effort. Suppose now that high effort individuals with bad circumstances have a tendency to migrate out of the measured population. The result will be that E and C are now positively correlated.

Suppose wages are set according to

$$w_i = \alpha C_i + \beta E_i.$$

Here α is the causal contribution of circumstances to wages, and β is the causal contribution of effort to wages. The structural model is graphically depicted in figure 3.

Like before, suppose E is not measurable and we instead estimate ψ by OLS on

$$w_i = \psi C_i + \varepsilon_i.$$

As shown in the previous example, the quantity of interest is the causal contribution α. However, we have

$$\mathbf{E}[\hat{\psi}] = \frac{\operatorname{cov}(w, C)}{\operatorname{var}(C)} = \alpha + \beta \frac{\operatorname{cov}(E, C)}{\operatorname{var}(C)}.$$[3]

Hence, assuming $\beta > 0$ and since E and C are positively correlated, the OLS estimator $\hat{\psi}$ overestimates the causal effect of circumstances on wages. (Similarly, if effort and circumstances are inversely correlated, $\hat{\psi}$ underestimates the causal effect of circumstances on wages.)

The problem can also be understood graphically. Since the data is gathered after migration has taken place, we are in effect conditioning on ('adjusting for') non-migration. In other words, sample selection S is part of the adjustment set Z. But Z does not satisfy the back-door criterion: S is a descendant of M, and M is a collider;

[3]This is a standard result in econometrics. See the omitted variable formula in Greene (2018, 59).

therefore, conditioning on S unblocks the path $C \rightarrow M \leftarrow E \rightarrow w$. This fails condition (ii) of the backdoor criterion.

In this example the regressor C is correlated with the error term ε. It should be noted that this is a type of bias that Bourguignon, Ferreira, and Menéndez (2007, 2013) attempt to correct for. They use a Monte Carlo method that considers a range of estimates of the bias of the estimated coefficients after drawing values for the correlation coefficients between regressors and the error term from a uniform distribution. It is claimed that this method is suitable when estimating a lower bound on the total amount of inequality of opportunity (Bourguignon, Ferreira, and Menéndez 2013).

4 How bad is the causal bias?

In the previous section I showed that the parametric inequality of opportunity measures proposed by Bourguignon, Ferreira, and Menéndez (2007) and Ferreira and Gignoux (2011) *may* be biased, but it could be maintained that the problem is small. I argue that this is not the case: the problem could be very large. As the extent of the problem cannot be measured, we are left in the dark about how accurate measures of inequality of opportunity really are.

Hitherto I have assumed that outcomes are completely determined by circumstances and effort (and on one occasion a random disturbance u). It is possible, however, that factors play a role that are classified as neither. For example, a common cause of something that is fully a circumstance and something that is fully a matter of responsibility cannot be classified. (If it is classified as a circumstance, say, then its effect on outcomes via the responsibility variable is a matter of luck, which means that the responsibility variable cannot be *fully* a matter of responsibility.) Moreover, it is possible to defend a version of equality of luck based on a negative principle about what people are not responsible for, without specifying what individuals *are* responsible for. The only requirement for equality of luck is that circumstances do *not* cause inequalities. Hence, I will now drop this assumption. Anything that is not a circumstance I will call a non-circumstance, which could or could not be a matter of responsibility.

As examples 3.1 and 3.3 show, there are two kinds of problems that bias the parametric *ex-ante* measure of inequality of opportunity. First, individuals' non-circumstances could have common causes with circumstances. Second, circumstances and non-circumstances could be correlated as a result of individuals leaving the measured population.

First consider the second problem. Deaths and migration occur all the time and should typically lead to correlations between circumstances and non-circumstances. This would be a serious problem unless (a) the changes are independent of circumstances and non-circumstances or (b) these changes in the population are very small. Option (a) is implausible, as it requires an unlikely lack of interaction between individual characteristics and circumstances. For example, adventurous people living in circumstances with low prospects may be more likely to migrate than adventurous people with better circumstances, so if adventurousness is a non-circumstance, (a) is not satisfied. Similarly, hard-working people living in a miner's community may be more likely to die. Or, ambitious people of wealthy families may be more or less likely to migrate to another country than poor ambitious people. On the other hand, (b) may be satisfied in some cases, such as when there is no migration and mortality in the measured population is sufficiently low. Hence, the second problem can potentially be overcome if researchers have good data. However, researcher should put effort into demonstrating that (b) is the case, which as far as I'm aware has never been done in the measurement literature.

Now consider the first problem, of common causes. Variables that (in the measurement literature) are typically considered circumstances are parental socioeconomic status, race, gender and locality. Non-circumstances, on the other hand, could be choice of education, occupation, and hours worked per week. If you would trace the causal history of these circumstances and non-circumstances, you would find many common causes. Individual characteristics such as interests and ambition are heavily influenced by parental upbringing, which in turns is associated with parental socioeconomic status. If talents are considered noncircumstances, then genetic characteristics are a common cause of talents and parental socioeconomic status. And so forth. In many such cases, however, some people would object along the lines that these common causes (such as genotype) are themselves circumstances. This would lead to a more expansive conception of a 'circumstance' that would, perhaps, make inequality of opportunity measures more accurate. Two versions of this objection will be discussed in section 4.1 and 4.2.

4.1 The free will objection

A special kind of objection is available to libertarians in the free will debate. According to the position I will call *cause-exclusive libertarianism*, individuals make free choices for which they are morally responsible only if they are entirely uncaused, and moreover, individuals are in fact capable of making free choices of this type. Suppose now that one restricts responsibility factors to such free choices, and one takes circumstances to be all things that are neither a free choice nor caused by a free choice. It follows that circumstances and non-circumstances have no common causes, eliminating any common cause bias. There might still be 'selection bias' of the type introduced in example 3.3, and inequality of opportunity might be underestimated due to unmeasured circumstances; but the problem would be much less severe.

However, cause-exclusive libertarianism is not an attractive position. First, the view that free actions must be entirely uncaused is controversial even among libertarians (Capes 2017). Most would agree that actions that have some causes but are not causally determined can be free choices of the kind that people are fully responsible for. For example, the choice of a poll worker to rig an election after being bribed is caused in part by the bribe offer, but the poll worker may still be fully responsible. Second, even if choices are only free if they are entirely uncaused, it is unlikely that individuals are capable of making such choices. Choices are always made within a context: in the context of growing up in a modern-era European city, someone will make very different choices than in the context of, say, 16th century rural China. Plausibly, all actions are in part caused by contextual factors. Consequently, free choices of the kind required in the argument do not exist, which means that a view of equality of opportunity based on it collapses to outcome egalitarianism.

4.2 The 'common causes are circumstances' objection

A couple of authors claim that if circumstances and effort are correlated, it is not ethically acceptable to hold people responsible for their effort (Fleurbaey 1998, 221; Checchi and Peragine 2010, 433). This proposition in itself might be defensible, but even if true, it is not obvious that one can then safely assume that circumstances and effort *are* uncorrelated. Without much justification, this is what these authors propose. (A similar point, not discussed here explicitly, is made by Roemer 2002, who argues that the outcome distribution of a type should be seen as a characteristic of the type.) I will give a more charitable interpretation of this objection. This version of the objection does not depend on libertarianism; it would be valid even if all outcomes, circumstances and non-circumstances are causally determined by earlier factors.

A causal (and more reasonable) version of the above proposition is that one is not fully responsible for an effort variable that has a common cause with a circumstance variable. Insofar as effort is caused by a circumstance or a cause of a circumstance, one is not responsible for this effort. In the context of the principle of equality of luck, this position would translate to: someone is not responsible for any causes of circumstances; causes of circumstances should themselves be considered circumstances. This still leaves room for non-circumstances, which would be those causes that are not caused by a circumstance and from which there is a causal path to the outcome that does not go via a circumstance.

More specifically, one could use the following procedure to select a particularly expansive set of circumstances, assuming that one has knowledge of the full causal mechanism behind outcomes. First, list a set of factors (causal variables) that one regards as clear circumstances. All other direct causes of the outcome of interest are initially listed as non-circumstances. Then, list all the common causes of circumstances and non-circumstances, and add these common causes to the list of circumstances. (Common causes of all characteristics of all individuals, such as 'the formation of planet earth' should be left out. One includes only causal variables that create differential outcomes.) Then, remove each non-circumstance that is now caused (in part) by a circumstance, but add as non-circumstance all causes of the removed non-circumstance that are not a circumstance or caused by a circumstance. The procedure is illustrated in figure 4. First, c_1 and c_2 are selected as clear circumstances and nc_1 remains as an initial non-circumstance. In figure 4a, f_5 is added as a circumstance, as it is a common cause of c_1 and nc_1. Then, nc_1 is removed as circumstance, and f_4 is added as a non-circumstance.

When this process has finished, any non-circumstances that are left have no common causes with circumstances. The procedure cannot be executed in practice, since we typically do not have full knowledge of the causal mechanism behind an outcome. However, with this conception of circumstances, common cause bias will be eliminated even if one only measures a subset of circumstances. Existing results that show inequality of opportunity measures can be seen as lower bounds if not all circumstances are measured (Ferreira and Gignoux 2011) are more plausible in the absence of common cause bias. For example, if the true causal mechanism is described by figure 4a and we measure only c_1 and c_2, one's inequality of opportunity estimate can plausibly be interpreted as a lower bound of the true inequality of opportunity.

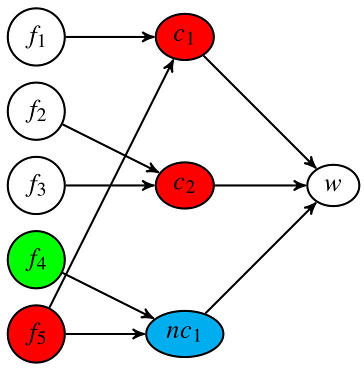

(a) Causal graph of w in which f_4 is a final non-circumstance.

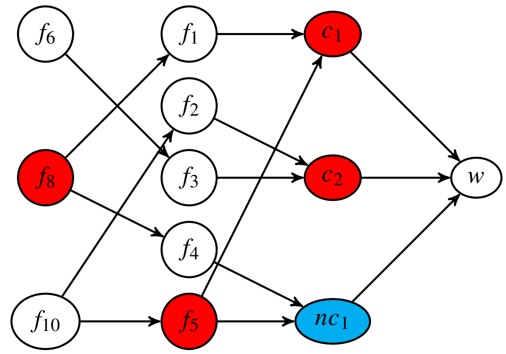

(b) Causal graph of w without final non-circumstances. All variation in w is caused by circumstances.

Figure 4: An illustration of the circumstance selection procedure. Circumstances are colored red. The initial direct causes of w that are non-circumstances are colored blue. The procedure leads to common causes of w and circumstances to be colored red. Final non-circumstances that remain are colored green.

However, there is a big danger associated with this procedure. It is likely that so many factors are selected as circumstances that no non-circumstances remain, or that their effect on outcomes is negligible. In that case, circumstances will be responsible for nearly 100% of inequality, such that one's position of inequality of opportunity collapses to outcome inequality. The problem is illustrated in figure 4b, which adds an additional layer of causes to figure 4a. In the more complex version, nc_1 now also has a common cause with c_1; namely, f_8. As a result, no non-circumstances remain.

This examples illustrates that when adding additional causal history to a graph, it becomes less likely that non-circumstances remain. Non-circumstances arise only if there are paths ending in the outcome variable of interest that never cross a circumstance variable. If the world is causally deterministic, this is an unlikely situation to arise for outcome variables (such as wages) that are typically measured, which have a very complicated causal history.

The following combinatorial argument makes this point more formally. Consider DAGs with a number of *periods*. In each period, there are q factors with arrows to two factors in the next period, with the requirement that each factor after the first period has at least one cause (it is allowed that both arrows point towards the same factor in the next period). In the last period (the most recent period before the outcome), there is one initial non-circumstance and there are $q - 1$ circumstances. It can be shown that the fraction of DAGs of this type which contain non-circumstances, out of all DAGs of this type, shrinks quickly towards zero when periods are added.[4] Hence, if reality is somewhat like this example, with a sufficient number of periods, it is unlikely that non-circumstances exist.

Perhaps an objector will find it more plausible than I do that non-circumstances

[4] The fraction of DAGs of this type, with $n + 1$ periods and q causes in each period, that contains non-circumstances is $(q - 1)^{n(q-1)}/q^{n(q-1)}$.

remain with our procedure. But a second problem with the objection is that it does not seem viable in a normative sense. As with cause-exclusive libertarianism discussed above, the procedure implies that individuals cannot be fully held responsible, in many situations, for choices that are in part caused by contextual factors. This should at least shift the burden of proof to the objector to show that the view is normatively acceptable. The categorization of circumstances is usually determined by reference to principles about choice and desert. The objector needs to show that the set of circumstances which the procedure generates is plausible given such principles.

Moreover, economists that measure inequality of opportunity have generally not used very expansive conceptions of circumstances, and often, like Roemer (1998), want to allow societies to decide for themselves what they consider circumstances and matters of responsibility. A society will not typically choose a set of circumstances that includes all the circumstances that ought be added if one were to follow our procedure. Hence, if the objector persist, she must concede that society does not have the last word on what circumstances are.

5 Measuring opportunity sets

This section discusses the non-parametric *ex-ante* approach, as it is used by Checchi and Peragine (2010) and Ferreira and Gignoux (2011), among others, and was introduced by Van de gaer (1993). The basic idea is to measure the value of each type's opportunity set by the arithmetic mean of that type's outcome distribution. Such an approach is normatively defended on the basis of equality of opportunity sets. This approach is called *utilitarian ex-ante*.

One imagines that each individual in society faces an opportunity set, which contains combinations of choices or other matters of responsibility (called *effort level*) and associated expected outcomes. It is assumed that two individuals of the same type face the same opportunity set. In the measurement literature it is also often assumed, as I do below, that all effort choices are available for all types.

Formally, this entails the following. The opportunity set C_t of type $t \in \{1, \ldots, n\}$ contains all combinations of possible effort levels $e \in \{1, \ldots, m\} = E$ and associated outcomes $f(e, t)$ that individuals of that type achieve with that effort level:

$$C_t = \{(e, f(e, t) \mid e \in \{1 \ldots, m\}\}.$$

We then need a valuation function in order to compare opportunity sets. An obvious choice is to sum all outcomes values in the set, leading to the valuation function $U(C_t)$.

$$U(C_t) = \sum_{e \in E} f(e, t).$$

(This valuation appears in Bossert 1997 and Ooghe, Schokkaert, and Van de gaer 2007. A different approach to valuing opportunity sets not discussed here is used by Lefranc, Pistolesi, and Trannoy 2008.)

Since effort is usually not measured, $U(C_t)$ is typically replaced by a slightly different function that can be estimated from data about circumstances only. A popular valuation due to Van de gaer (1993) identifies the value of a type's opportunity set with the mean income of that type, that is

$$\mu(t) = \frac{1}{N_t} \sum_{i \in t} x_i,$$

where N_t is the amount of people that are of type t and x_i is the outcome of individual i. We have assumed that $x_i = f(e_i, t_i)$, where e_i and t_i are the agent's effort level and type, respectively. Using $\mu(t)$ as valuation of opportunity sets leads to the utilitarian ex-ante approach for measuring inequality of opportunity.

Assuming that all circumstances are measured, $\mu(t)$ can be computed. Inequality of opportunity could then be defined as follows. We create a counterfactual distribution by replacing each individual's outcome with their type's mean $\mu(t)$, yielding

$$X^M = (\mu(1)1_{N_1}, \ldots, \mu(n)1_{N_n}).$$

Here N_t is the amount of people that are of type t and 1_{N_t} is a vector of 1's of length N_t. Inequality of opportunity can be estimated as

$$\tilde{\Theta}_I = I(X^M),$$

where I is some inequality index.

5.1 Problems with using mean income to value opportunity sets

The use of a type's mean income to value opportunity sets comes not without a cost (see also Hild and Voorhoeve 2004). Understood as a normative claim that the value of opportunity sets is given by $\mu(t)$ regardless of the underlying outcome distribution, it is deficient. This is because it is theoretically possible that two types face identical opportunity sets, but that individuals of one type just happen to choose different effort levels than individuals of the other type. This would lead to different estimations of the opportunity sets' value by $\mu(t)$, though by assumption, the opportunity sets are the same and must therefore have the same value. This shows that $\mu(t)$ is not universally acceptable, while it may still be empirically acceptable under suitable conditions, which I discuss below.

First, consider a simple example of a society in which there are two types with identical opportunity sets, depicted in table 1. Each individual has the option to choose either effort level e^1 with outcome 5 or effort level e^2 with outcome 10. However, 50 individuals of type 1 choose e^1, while 100 individuals of type 2 choose e^1. Moreover, 100 individuals from type 1 choose e^2, while 50 individuals of type 2 choose e^2. Assuming that the reason that individuals from type 2 choose effort levels with worse outcomes is innocuous, it is intuitively clear that there is equality of opportunity. This would also be our conclusion if we valued the opportunity sets using $U(C_t)$, which yields a value of 15 for both types. However, using $\mu(t)$ leads to a valuation of type 1's opportunity set of 8.34 and a valuation of type 2's opportunity set of 6.67.

The valuations U and μ do agree on there being equality of opportunity if both the opportunity sets of types are identical and effort is distributed identically across types.

	type 1	type 2
e^1	5 (50)	5 (100)
e^2	10 (100)	10 (50)

Table 1: Outcomes in a society with two effort levels and two types, with identical opportunity sets. The values in parentheses are the amount of individuals having the effort level and type of that cell.

	type A	type B
e^1	5 (100)	7 (100)
e^2	10 (50)	8 (50)

Table 2: A society with two types and two effort levels and identical effort distributions.

In what follows I make the case that the use of $\mu(t)$ for valuation is acceptable in general only if effort and circumstances are distributed independently.

Another way in which the (unconditional) distribution of effort affects $\mu(t)$ may not be so harmful, and we can use this fact to create a different valuation $U'(C_t)$ against which we can more easily compare $\mu(t)$. To use mean outcomes implies that effort levels which individuals are less likely to choose are given less weight. There is something to say for such a choice. For example, suppose opportunity set A offers a popular effort choice with very low rewards, and an unpopular effort choice with very high rewards. Opportunity set B, on the other hand, offers rewards that are only slightly below average for the popular effort choice, and rewards that are only slightly above average for the unpopular effort choice. It would make sense to value B above A, since more individuals would be benefited if they were given opportunities from B rather than A. This example is depicted in table 2. Using mean incomes as valuations concurs with this intuition: the mean income of type A is 6.67, while the mean income of type B is 7.33. On the other hand, we have both $U(A) = 15$ and $U(B) = 15$.

This consideration would favor the following valuation of a type's opportunity set. Let $P(e)$ be the probability that a given individual (regardless of type) has effort level e in the measured population. The valuation is given by

$$U'_P(C_t) = \sum_{e \in E} f(e,t)P(e). \tag{10}$$

This valuation scales an effort-outcome combination by the same factor $P(e)$ for each type, and is therefore not susceptible to the kind of normative problem discussed above with regards to the use of $\mu(t)$: two identical opportunity sets always have the same value, given a probability function P. If we apply this valuation to table 2, we get $U'_P(A) = 6.67$ and $U'_P(B) = 7.33$, the same values as the mean income assigns in this situation.

I submit that $U'_P(C_t)$ is a normatively acceptable valuation, and that $\mu(t)$ is empirically acceptable only insofar it agrees with $U'_P(C_t)$. That is to say, we require

$$\mathbf{E}[\mu(t) \mid t] = U'_P(C_t), \tag{11}$$

Here the expected value $\mathbf{E}[\mu(x_t) \mid t]$ is taken with respect to the probability function of effort, conditional on type. This assumption can of course be questioned, but someone who does so should explain which other valuation, if not $U'_P(C_t)$, ought to be tracked by $\mu(t)$, and I am aware of no such alternative. It seems $U'_P(C_t)$ comes closest to a valuation that (a) is normatively acceptable and (b) could reasonable inspire the use of $\mu(t)$ as empirical surrogate.

With this assumption, a condition under which $\mu(t)$ *is* empirically acceptable can be formulated. We can rewrite (11) as

$$\sum_{e \in E} f(e,t)P(e \mid t) = k \sum_{e \in E} f(e,t)P(e). \tag{12}$$

For equation (12) to hold in general for all t, we need that effort conditional on type is distributed identically for each type. (If not, the only other way in which (12) is satisfied is when different terms on the left side of the equation perfectly offset each other to equal their counterparts on the right, which will be rare.) In other words, effort and type should be independent. This we might call the assumption of randomness.[5]

> **The assumption of randomness.** For each type t_1 and t_2, the conditional distributions of effort are identical, that is, we have $P(e \mid t_1) = P(e \mid t_2)$.

This is assumption 2 in Checchi and Peragine (2010, 433), but it is typically left implicit in studies using an *ex-ante* approach.

By Reichenbach's common cause principle, the assumption of randomness is implied if there is no causal connection between effort and circumstances. If, to the contrary, there is a causal connection between effort and circumstances, either because one causes the other or if they have common causes, then the assumption of randomness is satisfied only under special conditions (see chapter 6 in Pearl 2009). Similarly, the assumption of randomness will fail to be satisfied if the population has changed (by death or migration) in a way that creates a correlation between effort and circumstances. Hence, the same sort of problems arise as with the parametric *ex-ante* approach. The arguments given in section 4 imply that the scope of the problem may be large for the non-parametric approach as well.

6 Conclusion

I argued that *ex-ante* approaches to measuring inequality of opportunity are biased due to causal confounding. This is the case both with parametric *ex-ante* methods, supposing that the type of inequality of opportunity they measure is inequality of luck, as well as non-parametric *ex-ante* methods, supposing that they measure inequality of opportunity sets.

Since *ex-post* methods do not necessarily depend on the measurement of causal quantities, my arguments do not extend to all *ex-post* methods. However, most *ex-post* studies depend on Roemer's Identification Assumption (RIA). This assumption states that an individual's effort level is her quantile in her type's outcome distribution (see also section 2). RIA depends on the asssumption of randomness (Fleurbaey 1998; Ramos and Van de gaer 2016). Hence, we should expect *ex-post* methods depending

[5]The name 'assumption of randomness' is inspired by what Sowell (1990) calls the "randomness assumption", the assumption that groups would be evenly represented by various outcomes measures, i.e. have identical outcome distributions, in the absence of unequal treatment of those groups.

on RIA, and by extension on the assumption of randomness, to be biased for similar reasons.

The problems discussed in this paper pertain well-known sources of bias in observational studies, particularly regression analysis. Such problems are most obviously tackled by switching to different methods, such as experimental methods utilizing randomization, or semi-experimental methods such as instrumental variables. The field of economics has moved in that direction in recent decades (Angrist and Pischke 2010). However, such methods are less readily available if the goal is to create a country-wide measure of inequality of opportunity. One way forward, which I favor, is to abandon the goal of creating reliable all-encompassing country-wide measures of inequality of opportunity. Instead, we could adopt the modest ambition to measure unfair contributions of particular circumstances in subgroups of the population. There is more hope that experimental and semi-experimental methods will be available for modest projects of that type.

References

Almås, Ingvild, Alexander W. Cappelen, Jo Thori Lind, Erik Ø. Sørensen, and Bertil Tungodden. 2011. "Measuring unfair (in)equality." *Journal of Public Economics* 95 (7): 488–499. doi:10.1016/j.jpubeco.2010.11.002.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2010. "The Credibility Revolution in Empirical Economics: How Better Research Design Is Taking the Con out of Econometrics." *Journal of Economic Perspectives* 24, no. 2 (June): 3–30. doi:10.1257/jep.24.2.3.

Arneson, Richard J. 1989. "Equality and equal opportunity for welfare." *Philosophical Studies* 56 (1): 77–93. doi:10.1007/BF00646210.

Balcázar, Carlos Felipe. 2015. "Lower bounds on inequality of opportunity and measurement error." *Economics Letters* 137:102–105. doi:10.1016/j.econlet.2015.10.026.

Bareinboim, Elias, and Judea Pearl. 2016. "Causal inference and the data-fusion problem." *Proceedings of the National Academy of Sciences* 113 (27): 7345–7352. doi:10.1073/pnas.1510507113.

Bareinboim, Elias, Jin Tian, and Judea Pearl. 2014. "Recovering from selection bias in causal and statistical inference." In *Twenty-Eighth AAAI Conference on Artificial Intelligence.*

Björklund, Anders, Markus Jäntti, and John E. Roemer. 2012. "Equality of opportunity and the distribution of long-run income in Sweden." *Social Choice and Welfare* 39 (2): 675–696. doi:10.1007/s00355-011-0609-3.

Bossert, Walter. 1997. "Opportunity sets and individual well-being." *Social Choice and Welfare* 14 (1): 97–112. doi:10.1007/s003550050054.

Bourguignon, François, Francisco H. G. Ferreira, and Marta Menéndez. 2007. "Inequality of Opportunity in Brazil." *Review of Income and Wealth* 53 (4): 585–618. doi:10.1111/j.1475-4991.2007.00247.x.

———. 2013. "Inequality of Opportunity in Brazil: A Corrigendum." *Review of Income and Wealth* 59 (3): 551–555. doi:10.1111/roiw.12045.

Brunori, Paolo, Flaviana Palmisano, and Vitorocco Peragine. 2019. "Inequality of opportunity in sub-Saharan Africa." *Applied Economics* 51 (60): 6428–6458. doi:10.1080/00036846.2019.1619018.

Capes, Justin A. 2017. "Freedom with Causation." *Erkenntnis* 82:327–338. doi:10.1007/s10670-016-9819-5.

Checchi, Daniele, and Vito Peragine. 2010. "Inequality of opportunity in Italy." *The Journal of Economic Inequality* 8 (4): 429–450. doi:10.1007/s10888-009-9118-3.

Checchi, Daniele, Vito Peragine, and Laura Serlenga. 2016. "Inequality of Opportunity in Europe: Is There a Role for Institutions?," 43:1–44. Research in Labor Economics. January. doi:10.1108/S0147-912120160000043008.

Cohen, G. A. 1989. "On the Currency of Egalitarian Justice." *Ethics* 99 (4): 906–944. doi:10.1086/293126.

Davillas, Apostolos, and Andrew M Jones. 2020. "Ex ante inequality of opportunity in health, decomposition and distributional analysis of biomarkers." *Journal of Health Economics* 69:102251. doi:10.1016/j.jhealeco.2019.102251.

Dworkin, Ronald. 2002. *Sovereign Virtue: The Theory and Practice of Equality.* Cambridge, MA: Harvard University Press.

Ferreira, Francisco H.G., and Vito Peragine. 2016. "Individual Responsibility and Equality of Opportunity." In *The Oxford Handbook of Well-Being and Public Policy,* ed. by Matthew D. Adler and Marc Fleurbaey. Oxford: Oxford University Press. Oxford Handbooks Online, doi:10.1093/oxfordhb/9780199325818.013.24.

Ferreira, Francisco HG, and Jérémie Gignoux. 2011. "The measurement of inequality of opportunity: Theory and an application to Latin America." *Review of Income and Wealth* 57 (4): 622–657. doi:https://doi.org/10.1111/j.1475-4991.2011.00467.x.

Fleurbaey, Marc. 1998. "Equality among responsible individuals." In *Freedom in Economics: New Perspectives in Normative Economics,* ed. by Nicolas Gravel Jean-François Laslier Marc Fleurbaey and Alain Trannoy, 206–234. London: Routledge.

Fleurbaey, Marc, and Vito Peragine. 2013. "Ex Ante Versus Ex Post Equality of Opportunity." *Economica* 80 (317): 118–130. doi:10.1111/j.1468-0335.2012.00941.x.

Greene, William H. 2018. *Econometric Analysis.* 8th ed. New York, NY: Pearson.

Hild, Matthias, and Alex Voorhoeve. 2004. "Equality of opportunity and opportunity dominance." *Economics and Philosophy* 20 (1): 117–145. doi:10.1017/S0266267104001282.

Hufe, Paul, Andreas Peichl, John Roemer, and Martin Ungerer. 2017. "Inequality of income acquisition: the role of childhood circumstances." *Social Choice and Welfare* 49, no. 3 (December): 499–544. doi:10.1007/s00355-017-1044-x.

Juárez, Florian Wendelspiess Chávez. 2015. "Measuring Inequality of Opportunity with Latent Variables." *Journal of Human Development and Capabilities* 16 (1): 106–121. doi:10.1080/19452829.2014.907247.

Lara Ibarra, Gabriel, and Adan L Martinez-Cruz. 2015. "Exploring the sources of downward bias in measuring inequality of opportunity." World Bank Policy Research Working Paper No. 7458, no. 7458. https://ssrn.com/abstract=2680470.

Lefranc, Arnaud, Nicolas Pistolesi, and Alain Trannoy. 2008. "Inequality of opportunities vs. inequality of outcomes: Are Western societies all alike?" *Review of Income and Wealth* 54 (4): 513–546. doi:10.1111/j.1475-4991.2008.00289.x.

Ooghe, Erwin, Erik Schokkaert, and Dirk Van de gaer. 2007. "Equality of opportunity versus equality of opportunity sets." *Social Choice and Welfare* 28 (2): 209–230. doi:10.1007/s00355-006-0165-4.

Pearl, Judea. 2009. *Causality.* 2nd ed. Cambridge: Cambridge University Press.

Pistolesi, Nicolas. 2008. "Inequality of opportunity in the land of opportunities, 1968–2001." *The Journal of Economic Inequality* 7, no. 4 (August): 411. doi:10.1007/s10888-008-9099-7.

Ramos, Xavier, and Dirk Van de gaer. 2016. "Approaches to Inequality of Opportunity: Principles, Measures and Evidence." *Journal of Economic Surveys* 30 (5): 855–883. doi:10.1111/joes.12121.

Roemer, John E. 1998. *Equality of Opportunity.* Cambridge, MA: Harvard University Press.

———. 2002. "Equality of opportunity: A progress report." *Social Choice and Welfare* 19 (2): 455–471. doi:10.1007/s003550100123.

Roemer, John E., and Alain Trannoy. 2015. "Equality of Opportunity." In *Handbook of Income Distribution,* ed. by Anthony B. Atkinson and François Bourguignon, 2:217–300. Oxford: North-Holland. doi:10.1016/B978-0-444-59428-0.00005-9.

Sowell, Thomas. 1990. *Preferential Policies: An International Perspective.* New York, NY: William Morrow / Company.

Van de gaer, Dirk. 1993. *Equality of opportunity and investment in human capital.* Leuven: KULeuven. PhD thesis.