

Age Discrimination: Is It Special? Is It Wrong?

Abstract

In this paper, I examine the moral status of age discrimination, by bringing together accounts of the wrongness of discrimination in general, and of the specialness of age discrimination in particular. I summarise recent attempts to account for the specialness of age discrimination and develop a taxonomy for different forms of age discrimination. I then relate three influential accounts of the wrongness of discrimination to this taxonomy. I show that different forms of age discrimination are classified as wrongful by different accounts. I moreover explore the roles that considerations of the specialness of age can play under these different accounts.

1 Introduction

Imagine Daryl and Eliot, two qualified applicants for a currently vacant position. None of them are called to be interviewed for the job – simply because of their gender (Daryl is female) and race (Eliot is black), respectively. As described, these cases are intuitively clear instances of discrimination, and as such intuitively morally wrong.

They are classified as discrimination also by the following definition of group discrimination:

- (GD) An agent, X, group discriminates against someone, Y, by ϕ -ing if, and only if:
- (i) There is a property, P, such that Y has P or X believes that Y has P,
 - (ii) By ϕ -ing, X treats Y worse than X would have treated Y, had Y not had P or had X not believed Y to have P,
 - (iii) It is because (X believes that) Y has P that X treats Y worse by ϕ -ing, and
 - (iv) P is the property of being a member of a socially salient group [i.e., a group, perceived membership of which is important to the structure of social interactions across a wide range of social contexts].¹

Being female and being black, respectively, are socially salient properties in virtually all societies. Clearly, not being called to be interviewed for a job that one applied and is qualified for because of such a property – when one would have been called had one lacked the property – amounts to worse treatment in at least one of the following senses:

- (a) being made worse off (e.g., by being deprived of an opportunity) or
- (b) being treated as inferior (e.g., by being considered not worthy of equal consideration with other – male or white – applicants).²

And arguably, treating someone worse in at least one of these senses, just because of their socially salient property, is *prima facie* morally wrong. Hence, *ceteris paribus*, Daryl and Eliot are wrongfully discriminated against.

Now imagine another case: three applicants, Alex, Billie and Charlie, are equally qualified for a currently vacant position. They belong to three different birth cohorts (e.g., three consecutive generations) and thus to different age groups: Alex is 60+, Billie is 40 and Charlie

¹ (Berndt Rasmussen 2018, sec. 7). For similar definitions, see (Lippert-Rasmussen 2014) and many of the entries in (Lippert-Rasmussen 2017).

² Cf. (Berndt Rasmussen 2018, sec. 7).

is under 18 years old. While one of them, Billie, is called to be interviewed, the other two are not – simply because of their old and young age, respectively.

Clearly, being old (60+) and being young (under 18) are socially salient properties in virtually all societies, and in this sense comparable to being female and being black. Moreover, in not being called to the interview, Alex and Charlie are exposed to the same form of worse treatment as Daryl and Eliot. Hence, even their cases should constitute clear-cut discrimination, according to the above definition (and, surely, according to many people's judgment at first sight). And hence, their treatment seems equally morally wrongful as Daryl's and Eliot's, and for the same reasons.

However, writers on age discrimination have pondered the thought that age is "special" – that age discrimination is relevantly different from other forms, e.g., gender and racial discrimination. Such specialness in turn might have moral ramifications, possibly making age discrimination less severe, or more justifiable, than these other forms. Common practices and intuitions seem to support this idea: in many societies, it is common and commonly accepted that, e.g., the right to vote and to run for office is not granted to minors; that car or life insurance premiums are age-adjusted; or that retirement at a certain age becomes mandatory. On the other hand, we might ask: could there be good reasons to change these practices and disregard these intuitions instead?

This paper examines the moral status of age discrimination, by bringing together accounts of the wrongness of discrimination in general, and of the specialness of age discrimination in particular. Section 2 summarises recent attempts to account for the specialness of age discrimination. In section 3, I develop a taxonomy for different forms of age discrimination. In section 4, I relate three influential accounts of the wrongness of discrimination to this taxonomy. I show that different forms of age discrimination are covered by different accounts. I moreover return to the specialness considerations and explore the roles they can play under these different accounts.

2 The specialness of age and age discrimination

The debate on age discrimination largely focuses on chronological age, defined as the number of years from a person's birth to the given date. Chronological age (I will henceforth omit the qualifier) has at least three special features – as compared to e.g. gender or race (as typically construed).³ It is:

- a) *non-discrete*, such that a normal life span can be partitioned into any arbitrary number of age segments;⁴
- b) *transient*, such that, for each person with a normal life span, they will belong to a number of such age segments as time passes; and
- c) perfectly *passage-of-time correlated*, due to the above definition.

A key idea is that, because of these special features, relying on age criteria (such as minimum age requirements) may sometimes be justified by reference to *efficiency* in two

³ Cf. (Bidadanure 2016, 247), (Macnicol 2006, 25), (Gosseries 2014), (Cupit 2013).

⁴ (Cupit 2013) however cautions that the arbitrariness in partitioning may also lead to arbitrariness in treatment.

different senses. (1) Age is a reliable proxy (statistical indicator) for certain given target variables and as such a potential tool for more efficient decision making. (2) Using age criteria leads to better overall outcomes, in some sense.

2.1 Age as a reliable proxy

In typical employment cases, applicant age is “significantly correlated”⁵ to job qualifications and expected productivity, due to specialness features (b) transience and (c) passage-of-time correlation. To illustrate: an applicant under 18 (such as Charlie) most likely does not have a college or university degree. This has to do with the way our societies organise their education systems (age of entry, sequence of educational levels, number of years required for a degree, etc.). Likewise, such an applicant is likely not to have the cognitive capacities required for e.g. jobs where risk management is central. This has to do with typical neurological/brain development. An applicant over 60 (such as Alex) is likely to lack some of the physical and cognitive abilities of an applicant in their 40’s (such as Billie), but also likely to have more experience-based competences. This has to do with processes of physical and mental decline, and with the time dimension of learning curves, respectively.

Likewise, in a medical context, given statistical facts about life expectancy, age can work as a reasonable proxy for remaining life years.⁶ Thus, age can be relevant for allocating scarce health resources, usually in combination with other factors, as exemplified by the econometric concepts of quality-adjusted life year, QALY, and disability-adjusted life year, DALY.⁷

Thus, for a wide variety of reasons, age can work as a reliable proxy for relevant individual target variables. Of course, for any given applicant or patient, an employer or medical provider might make a more accurate assessment of such target variables by conducting a (typically more time-consuming or complicated) individual assessment. Yet, to make their decision-making more efficient, they may choose to rely on age criteria instead. Thus, in promoting *procedural efficiency*, age may play a special role.

A further attempt to justify reliance on age criteria derives from the special feature of non-discreteness, (a): that a normal lifespan is divisible into any set of age segments that are relevant in a given context. This means that using age criteria potentially allows for higher levels of “precision”⁸ in approximating the relevant target variables, compared to the (typically) discrete properties of gender and race. This can further promote procedural efficiency by making age a more reliable proxy.

There is a general objection against this reliable-proxy justification of age criteria.⁹ In openly sexist or racist societies, where women or people of colour are denied higher education, gender and race would also be reliable – and sufficiently precise – proxies. Still, we would not consider the resulting discriminatory hiring decisions as less severe, or more justifiable than their counterparts in our (less openly sexist or racist) societies, due to these different

⁵ Cf. (Gosseries 2014, 63f.).

⁶ Cf. (Cupit 2013), (Bognar 2015).

⁷ For an overview of the debate on the ethics of health care allocation, see (Bickenbach 2016).

⁸ (Gosseries 2014, 63f.), (Bidadanure 2016).

⁹ Cf. (Gosseries 2014, 65f.).

social facts. Whether something is a reliable proxy is contingent on e.g. physical, biological, and psychological facts, but also social facts that depend on (the aggregate of) our choices, e.g. on how we choose to set up our education system. When social facts are invoked to justify the use of a proxy, the justificatory burden shifts to these social facts and underlying choices. That is, if we seek to justify the use of age criteria by appealing to the fact that age is a reliable proxy, yet this latter fact is explained by the fact that our society is organised around age criteria, our argument becomes circular unless we can independently justify society's organisation around age criteria.

This line of thought thus leads us to consider the second sense of the specialness of age.

2.2 Age criteria improve overall outcomes

This second efficiency concern is about *outcome efficiency*, and as such distinct from procedural efficiency. That is to say, we could reject the idea that age is a reliable, sufficiently precise proxy for specific target variables, and still concede that relying on age criteria allows for better overall outcomes. 'Better' may here be understood in terms of overall utility or in terms of fairness or equality.

2.2.1 Utility

One key account under this headline is that of societal "sequence efficiency":¹⁰ organising society around age criteria, in certain contexts, promotes total utility, as it makes society overall more functional. The education system is an obvious case in point. There are efficiency reasons for making people go through consecutive stages of education, prior to taking higher levels, and prior to entering the labour market. Again, however, one might object that there might be overall efficiency gains from race or gender discrimination, e.g. in a caste-like hierarchy, which we would not grant justificatory force.

With age, though, there is a crucial difference. As the education system example shows, *societal* sequence efficiency here originates from *individual* sequence efficiency.¹¹ In certain contexts, ordering someone's activities in chronological sequences promotes her achievement of the context-given objectives, due to specialness feature (c): perfect correlation with the passage of time, which in turn correlates with other factors. E.g., learning how to spell prior to signing up for creative writing classes yields a better outcome for the individual than the reverse. Individual sequence efficiency is thus a further account of the specialness of age – and one with greater justificatory force, since the efficiency trade-offs it is based on are intra-personal rather than interpersonal (as in the sexist or racist caste system).

2.2.2 Fairness/equality

From a perspective of fairness or equality, age criteria are problematic as they imply worse treatment (in at least one of the above senses) of some individuals at some specific time,

¹⁰ (Gosseries 2014, 70f.).

¹¹ (Bidadanure 2016, 247) calls this "lifespan efficiency": for any given individual, certain patterns of resource distribution across different age segments maximise utility across her entire lifespan. Yet another way to address this issue is to ask which resource allocation pattern over her lifetime a rational chooser would prefer from behind a Rawlsian veil of ignorance (Cupit 2013); cf. (Daniels 1988), (Bognar 2015), (Bognar 2008).

due to a socially salient feature. Young Alex and old Charlie are denied an employment opportunity, which is granted to middle-aged Billie. The key idea with justifying such differential treatment, nevertheless, from a fairness or equality perspective is that, in the case of age, such inequality in treatment is mitigated, or indeed eliminated, once we see the whole picture.

The most ambitious attempt of justification is the “complete life neutrality” account:¹² given some set of age criteria, *ceteris paribus*, there will be no inequality in treatment once we consider complete lives. The *ceteris paribus* clause ensures that – due to specialness feature (b), transience – each age criterion for resource allocation applies equally to everyone, if only at different points in time, thus treating each equally over their lifetime.

However, this presupposes an overly strong *ceteris paribus* condition, where everyone has the same life length as well as the same lifetime profile of needs, goals and desires. E.g., a medical treatment offered to all under 70 affects those who happen to need it at 65 or 75, respectively, very differently; a pensions plan for all and only those above a certain age benefits only whoever lives that long.¹³ Moreover, any changes in such age criteria – or in the environment in which these criteria receive significance – may lead to complete life inequalities between birth cohorts. E.g., if the required age for a benefit is raised at some point in time, all birth cohorts reaching the threshold age after that point will be worse off, over their lifetime, than previous ones.

Moreover, one could object that even if complete life neutrality could be achieved, inequalities that exist at any specific time between age groups matter in themselves, and they do not disappear because of “reversed” inequalities at some other time. The idea is that relational inequalities between people (or groups) cannot be compensated in the same way as one person’s (or group’s) bad states at some time can be compensated by the same person’s (or group’s) good states at some other time. One possible explanation for this is that inequality simply matters in itself.¹⁴ Another is a time-specific prioritarian view, where the value of a benefit for someone at some point in time depends on this person’s relative welfare level at this point.¹⁵

An attempt to mitigate fairness or equality-based reasons against age discrimination comes from the “diversification” account:¹⁶ even if age criteria result in complete life inequalities, the latter are moderated due to age’s specialness. Due to feature (a), non-discreteness, actual age limits tend to vary across different contexts – e.g. 14, 16, 18, 21 – perhaps somewhat arbitrarily or because they are finetuned to different purposes. Hence, resulting disadvantages will spread out more evenly over an individual’s lifetime.

¹² (Gosseries 2014, 66ff).

¹³ Cf. (Cupit 2013). (McKerlie 1992) suggests that, due to differences in life span, we should discriminate in favour of the young, to concentrate resources to life-stages through which more will live.

¹⁴ See Bidadanure’s relational conception of equality (which does not replace but complements the distributive conception of equality over complete lives), (Bidadanure 2016).

¹⁵ Cf. (McKerlie 1992), (McKerlie 2001).

¹⁶ (Gosseries 2014, 63).

There is a further argument that relying on age criteria allows for better overall outcomes. According to an “affirmative egalitarian” account,¹⁷ this may actively reduce given social inequalities over complete lives. Consider e.g. mandatory retirement at age 67. Assume that there is a number of consecutive birth cohorts, each of which has a certain proportion of unemployed individuals. As one cohort reaches the threshold, their positions are made vacant and filled with members of the subsequent cohorts, shifting employment benefits to those previously unemployed. Later on, these cohorts of course have to do the same for their successors. This means that people’s life trajectories are equalised in the sense that fewer will go through stretches of unemployment *prior to*, while all are retired *after*, the set age threshold.¹⁸ Another example is the allocation of life-saving treatment: fairness requires that we give it to a young person who has not yet reached the “fair innings” threshold of a “complete or full life”, rather than to an old person who is beyond this threshold.¹⁹

2.3 The argument from specialness

The general argument from specialness considerations can then be summarised as follows:

- (1) There is a *prima facie* reason *against* discrimination as defined by (GD).
- (2) Using age criteria is a form of discrimination (just as using gender/race criteria).
- (3) Specialness: in some cases or contexts C, there are efficiency reasons (of the procedural and/or outcome variety) *for* using age criteria, which outweigh the *prima facie* reasons against it.
- (4) Hence, in C, we have overall reason *for* using age criteria.

The above efficiency considerations thus constitute different ways of precisifying premise (3).

3 A taxonomy of age discrimination

The above procedural efficiency accounts of the specialness of age may serve to justify cases where age is used as a reliable proxy in order to approximate individual target variables. Whenever an agent’s use of age criteria is based on a statistical correlation between age and the relevant target variable, we are dealing with *statistical discrimination*. We can capture this phenomenon by precisifying condition (iii) of the above definition (GD):

(iii*) It is because X believes that Y has P, *and that P is statistically correlated to a target variable that is of relevance for X*, that X treats Y worse by ϕ -ing.²⁰

The justificatory force of procedural efficiency considerations for using age criteria then arguably hinges on (1) whether the agent’s belief, that P is statistically correlated to the relevant target variable, is correct – as well as (2) whether the goals the agent pursues through these target variables can be justified.

¹⁷ (Gosseries 2014, 70ff).

¹⁸ Cf. (Wedeking 1990). For a dissenting view, see e.g. (Overall 2006), who decries mandatory retirement as a “common socially disabling practice”; cf. (Nussbaum and Levmore 2017).

¹⁹ For critical discussion, see (Bognar 2008).

²⁰ Cf. (Lippert-Rasmussen 2014, 81). In Schauer’s words: a form of discrimination that is “alleged by the discriminator to rest on a statistically valid empirical distinction – a correlation – between two (or more) people or categories” (Schauer 2017, 42f.).

The outcome efficiency accounts may help justify cases where age criteria are a means to overall better outcomes, even if age is not considered a reliable proxy for any specific individual target variables. Even here, a society (as a collective agent X) that stipulates age limits could be said to statistically discriminate against certain age groups (Y), in the following, wider sense. If public attempts to *justify* such age limits refer to statistical evidence relating the general use of such criteria to certain desired social outcomes, society can be said to engage in differential treatment of different age groups *because of* such statistical correlations. E.g., if the justification of certain age limits on alcohol consumption appeals to epidemiological evidence of their correlation to reduced youth criminality or disease rates, such criteria are used not because they are proxies for individual target variables, but because they are means for achieving a better societal outcome.

Statistical discrimination is typically considered as an instance of direct discrimination.²¹ Direct discrimination is usually taken to mean *disparate treatment* due to the agent's *intention* to treat people differently on the basis of their membership of a socially salient group.²² And this is certainly the case when an agent intentionally uses age as a proxy – or when a collective agent's use of age is due to the idea that this is a means for improving the social outcome. In contrast, indirect discrimination is typically understood to involve no intentions to treat differently – while treating differently nonetheless.²³ This is the case when a “facially neutral” policy *unintentionally* has a *disparate impact* on members of different age groups. As Macnicol states:

Direct age discrimination is openly stated, and obvious to all (such as age limits in job advertisements). Indirect discrimination is that which is apparently neutral and therefore not openly admitted, or not necessarily age-specific, but which has an adverse impact upon older people. [...] In the case of indirect discrimination, [disparate] outcome matters more than intent: a ‘facially neutral’ policy can have an adverse impact on a particular group, even if no intent to discriminate exists.²⁴

This common way of drawing the distinction, however, brings out that there are two separate distinctions at work.²⁵ One is between disparate treatment – ϕ -ing or χ -ing towards people with or without P – and disparate impact – affecting people with or without P

²¹ Cf. (Lippert-Rasmussen 2014, 88). Lippert-Rasmussen suggests there can be indirect statistical discrimination, as when a city council decides to let their police force use race as a proxy for differential law enforcement measures, where the city council is ignorant of how this will effect different racial groups, i.e. which group will turn out to be at the worse end of the resulting measures. However, this example won't work: the council's decision (“you may use racial profiling”) is in itself not discriminating against anyone. Yet any subsequent race-based police measure (“we'll stop-and-search primarily foreign-looking people”) will be *direct* discrimination on the basis of race.

²² Cf. (Lippert-Rasmussen 2017, 3).

²³ Cf. (Lippert-Rasmussen 2017, 3). (Altman 2016, §3.1) states: “Direct discrimination is essentially a matter of the reasons or motives that guide the act or policy of a particular agent, while indirect discrimination is not about such reasons or motives.” Cf. even (Moreau 2017, 166–67).

²⁴ (Macnicol 2006, 21).

²⁵ Cf. (Berndt Rasmussen (ms.)).

differently by ϕ -ing. This can be brought out by two alternative precisifications of condition (ii) in GD:

(ii') Disparate treatment: Had Y not had P, X would have χ -ed, rather than ϕ -ed, and ϕ -ing toward someone constitutes worse treatment of them than χ -ing.

(ii'') Disparate impact: Had Y not had P, X would still have ϕ -ed, but ϕ -ing toward someone with P constitutes worse treatment than ϕ -ing toward someone without P.

The other distinction is between intentional and non-intentional treatment and concerns P's place in the action's explanation: as part of X's motivating reasons, or of some other explanatory reason. This can be brought out by two alternative precisifications of condition (iii):

(iii') Intentional: It is because X has P-related intentions – e.g. *dislikes* people with P and *believes* that Y has P – that X treats Y worse.

(iii'') Non-intentional: It is *not* because X has P-related intentions that X treats Y worse, but rather because of some other reason.

We thus have two separate distinctions, each of which has been used – though in a conflated fashion – to define direct and indirect discrimination. My contention is that both the disparate treatment of different groups, and the disparate impact on different groups of the same, “facially neutral” treatment, can be either intentionally or non-intentionally brought about by the agent.²⁶ Thus, we are really dealing with four, rather than two different types of discrimination (see *Table 1*).

	<i>Disparate treatment: ϕ-ing rather than. χ-ing</i>	<i>Disparate impact of ϕ-ing</i>
<i>Intentional</i>	(I) Intentionally using age as a criterion for differential treatment, e.g. setting a hiring age of 25/a minimum employment age of 14.	(III) Employing a facially neutral criterion, e.g. “overqualified”, while intending its differential impact on older applicants.
<i>Non-intentional</i>	(II) Treating differently without intention, e.g. perceiving equally qualified older applicants as less qualified, due to implicit age bias ²⁷	(IV) Employing a facially neutral criterion, e.g. “overqualified”, without intending its differential impact on older applicants.

Table 1

Age discrimination, understood as statistical discrimination according to the above definition with precisified condition (iii*), can fall under either disparate treatment or disparate impact on the intentional dimension, i.e. (I) or (III). E.g., an employer who sets a minimum hiring age of 25 due to physiological evidence that under-25-year olds tend to have underdeveloped brain regions for risk assessment, statistically discriminates against under 25-year olds. Likewise, an employer whose use of the criterion “overqualified” is explained by a preference for younger employees can be said to statistically discriminate against older

²⁶ See (Berndt Rasmussen (ms)).

²⁷ Cf. For empirical studies of implicit age bias, see e.g. (Malinen and Johnston 2013), (Cortina et al. 2013), (Derous and Decoster 2017).

applicants, if this preference is due to her belief that younger age correlates with “more grit”. (If this preference instead is due to e.g. her fear of aging, she does not statistically discriminate against older applicants.) A society which sets the minimum employment age to 14 and justifies this with e.g. empirical evidence that such a measure tends to increase the education level of populations can be said to engage in statistical discrimination of the wider kind. (If this policy is instead justified by reference to religious doctrine, this is not statistical discrimination.) For each such instance of statistical discrimination, the question then arises whether the specialness of age generates efficiency reasons which outweigh the prima facie reason against discriminating.

As these examples suggest, there are forms of intentional age discrimination which do not fall under the label of statistical discrimination. And as *Table 1* suggests, there are forms that do not constitute any intentional form of discrimination. What is the role of specialness considerations in these cases? They do not form part of the agent’s intentions or society’s wider justification attempts. However, might they nevertheless justify reliance on age discrimination, in certain cases or contexts C? How should we weigh efficiency reasons, according to the different specialness accounts, against the prima facie reasons against discrimination, which may rely on a variety of moral frameworks? In the final section, I will consider three prominent accounts of the wrongness of discrimination and explore how they connect to the issue of age discrimination and the specialness of age.

4 The wrongness of age discrimination

4.1 Mental state accounts

According to one influential mental state account, an instance of discrimination is prima facie wrong when and because it is “premised on the belief that some types of people are morally worthier than others”, in the sense of meriting greater moral concern.²⁸ So if e.g. Alex, who is 60+, is not called back for an interview and this is due to the employer’s belief that older applicants are just not worth the trouble (or due to Alex failing some other facially neutral criterion, posited because of such a belief), this constitutes wrongful age discrimination. If, on the other hand, the employer’s decision is due to an implicit (unconscious, unendorsed) age bias, the account fails to classify it as morally wrong. According to another mental state account, an instance of discrimination is prima facie wrong when and because it is “done from unwarranted animus or prejudice against persons of [a socially salient] type”, where “animus” means “a negative attitude, an aversion” and “prejudice” a belief with “faulty origin”.²⁹

In so far as these accounts focus on the beliefs and/or attitudes that constitute the motivating reasons behind the act in question, they can generate prima facie reasons against both types of intentional discrimination – as long as the motivating beliefs or attitudes concern moral inferiority, or are unwarranted or of faulty origin, respectively. When, on the other hand, these motivating beliefs and attitudes concern appropriately grounded efficiency considerations, e.g. when age is intentionally used as a proxy, this is not morally wrong. Statistical discrimination, properly informed, thus gets a moral free pass.³⁰

²⁸ (Alexander 1992).

²⁹ (Arneson 2006, 787f.).

³⁰ Cf. (Alexander 1992).

Specialness considerations, if they are part of the mental content that motivates a discriminating act, can mitigate the moral status of the act in question.

On the other hand, since these mental state accounts focus on motivating reasons, i.e. beliefs and/or desires, they cannot classify non-intentional discrimination as morally wrong. There is then no *prima facie* reason against e.g. differential treatment from implicit age bias.

4.2 Social meaning accounts

According to social meaning accounts, an instance of discrimination is *prima facie* wrong when and because it is objectively *demeaning*, in the sense of expressing “a lack of respect for the equal humanity of the other” from a position of power that gives force to the expression.³¹ To assess whether these conditions are met, we need to make an interpretative judgment of the act, and this in turn requires us to examine the social context: the conventional methods of showing respect and disrespect, the status of the agent, etc. E.g., the decision to not hire Alex who is 60+ is objectively demeaning if it is delivered, from the authoritative position of an employer, in a way that patronises Alex or taps into stereotypes of elderly people. An evaluation of such factors can of course be neither entirely independent of people’s beliefs and practices (strongly objective), nor totally dependent on what the majority thinks (minimally objective). Thus, demeaningness is at most a *modestly* objective feature of an act.³²

Note that, typically, on social meaning accounts, what matters is a relational, synchronic inequality between people: discrimination is wrong when and because it treats “people in a way that is demeaning to their synchronic persons”: it denies another “her equal moral worth” at a given point in time.³³ This can have repercussions for the argument from specialness. Suppose we want premise (3) to appeal to outweighing efficiency reasons from equality under the complete life neutrality account. This account asserts what a (synchronic, relational) social meaning account denies: that inequalities between people at one point in time can be compensated by “reversed” inequalities at another. This means that on a social meaning account reading of premise (1), we cannot plausibly appeal to a complete life neutrality account on premise (3), without further qualification.

Since social meaning accounts focus solely on what discriminating acts express – regardless of their underlying intentions, and regardless of whether they constitute differential treatment or have differential impact – such accounts have the potential to cover all the above four types of discrimination. However, for each instance of discrimination, its moral assessment will require interpretative judgments. Though there is likely to be agreement about the social meaning of some instances, there will undoubtedly be disagreement about the meaning of others. This leads some to worry that whenever there is pervasive disagreement – as there tends to be on socially sensitive issues like group discrimination – there will be no (sufficiently) objective social meaning to instances of differential treatment.

However, the potential for contestation on this account of the wrongness of discrimination may be a feature, not a bug: while the account does not give ready-made answers, it

³¹ (Hellman 2008, 36). Cf. (Bidadanure 2016), (Shin 2009).

³² (Hellman 2008, 75).

³³ (Hellman 2008, 29).

provides a template for collective deliberation concerning the moral status of differential treatment along socially salient lines. In the case of age discrimination, this is where specialness considerations can play an additional role: they can help rewrite the social meaning of acts within such collective deliberation. E.g., if some agreement is reached that using a minimum employment age is socially beneficial, then a decision to not hire a 13-year old, just because of their age, will not be objectively demeaning (in the moderately objective sense), and hence not morally wrong to start with.

Specialness considerations will thus not only have the potential to generate outweighing efficiency reasons (premise (3)) but also to directly modify the prima facie reasons against age discrimination appealed to in premise (1).

4.3 Harm accounts

While the above accounts of the wrongness of discrimination pick up on one sense of the 'being treated worse' clause in (GD) – (b) being treated as inferior – the present account picks up on the other sense – (a) being made worse off. According to harm accounts, an instance of discrimination is prima facie wrong when and because it harms the discriminatee, i.e., makes them worse off as compared to what they would have been, had they not had socially salient property P.³⁴ Since such harm accounts focus solely on the comparative outcomes of discriminating acts – regardless of their underlying intentions, and regardless of whether they constitute differential treatment or have differential impact – they have the potential to cover all the above four types of discrimination.

Consider Alex who is 60+ and is not called back for an interview. Had she been middle-aged, she would have been called back (as the evidence from Billie's case indicates). Has she been made worse off compared to a counterfactual state in which she had been middle-aged, rather than 60+? This may seem obvious at first. Yet in the light of the above complete life neutrality account, there is room for doubt. There are compelling reasons why the morally significant units to be compared here are complete lives.³⁵ And considering her complete life, *ceteris paribus*, Alex is not made worse off: had she been middle-aged at this point in time (like Billie), she would have been 60+ at some later point, and hence not been called for an interview *then* instead. Maybe we could say that Alex is made worse off, compared to had she been middle-aged (like Billie) at all and any times of her life. Yet this is nomologically impossible, and hence cannot be a relevant counterfactual here. There is still another sense in which we could claim that Alex is made worse off, viz., compared to the counterfactual state in which no age criteria are used at all. Assessing this claim is a difficult empirical matter. Specialness considerations pertaining to individual (and societal) sequence efficiency give us at least some reasons to doubt whether such a claim would eventually be born out.

However, there may be circumstance where the strong *ceteris paribus* clause does not hold (e.g. since employment opportunities vary over different birth cohorts), or where age criteria do not achieve social or individual sequence efficiency. Then, it may be true that an individual like Alex is made worse off by the use of age criteria than she would have been in the relevant counterfactual state, and that there is, thus, a prima facie reason against treating her differently due to her age. Then, again, specialness considerations – e.g. to the

³⁴ This is a counterfactual, welfarist sense of harm, cf. (Berndt Rasmussen 2018), (Lippert-Rasmussen 2014).

³⁵ (Bidadanure 2016, 244f.)

effect that age criteria such as mandatory retirement rules reduce inequalities (on the above affirmative egalitarian account) – may be invoked as reasons which outweigh these prima facie reasons.

5 Conclusions

As initially stated, the use of age criteria is both common and commonly accepted in many societies. Does using age criteria constitute age discrimination, and does this in turn give us good reasons against such a practice?

First, and rather obviously, an answer to this question depends on which account of the wrongness of discrimination we settle for. Second, we need to be clear about which type of age discrimination we are considering. I suggested that there are at least four types, drawn up along the double distinction of intentional/non-intentional disparate treatment/disparate impact discrimination, where statistical age discrimination falls under one or the other type of intentional discrimination. The different accounts of the wrongness of discrimination classify different types as prima facie morally wrong. Third, the different accounts of the wrongness of discrimination are sensitive in varying extents to accounts of the specialness of age, which in turn are due to the three special features of age: non-discreteness, transience and passage-of-time correlation. Fourth, whenever there is a prima facie reason against using age criteria, according to the relevant account of the wrongness of discrimination, the argument from specialness provides room – but also certain restrictions – for procedural or outcome efficiency considerations to generate outweighing reasons in favour of using age criteria.

These conclusions do not settle the matter of the moral status of using age criteria across the board. Hopefully, though, they may serve as a template for structuring and morally evaluating specific instances of age discrimination.

Bibliography

- Alexander, Larry. 1992. 'What Makes Wrongful Discrimination Wrong Biases, Preferences, Stereotypes, and Proxies'. *University of Pennsylvania Law Review* 141 (1): 149.
- Altman, Andrew. 2016. 'Discrimination'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2016. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/win2016/entries/discrimination/>.
- Arneson, Richard. 2006. 'What Is Wrongful Discrimination?' *San Diego Law Review* 43 (4): 775.
- Berndt Rasmussen, Katharina. (ms.). 'Implicit Bias and Discrimination'.
- . 2018. 'Harm and Discrimination'. *Ethical Theory and Moral Practice*, August. <https://doi.org/10.1007/s10677-018-9908-4>.
- Bickenbach, Jerome. 2016. 'Disability and Health Care Rationing'. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2016. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2016/entries/disability-care-rationing/>.
- Bidadanure, Juliana. 2016. 'Making Sense of Age-Group Justice: A Time for Relational Equality?' *Politics, Philosophy & Economics* 15 (3): 234–60. <https://doi.org/10.1177/1470594X16650542>.

- Bognar, Greg. 2008. 'Age-Weighting'. *Economics & Philosophy* 24 (2): 167–89. <https://doi.org/10.1017/S026626710800179X>.
- . 2015. 'Fair Innings'. *Bioethics* 29 (4): 251–61. <https://doi.org/10.1111/bioe.12101>.
- Cortina, Lilia M., Dana Kabat-Farr, Emily A. Leskinen, Marisela Huerta, and Vicki J. Magley. 2013. 'Selective Incivility as Modern Discrimination in Organizations: Evidence and Impact'. *Journal of Management* 39 (6): 1579–1605. <https://doi.org/10.1177/0149206311418835>.
- Cupit, Geoffrey. 2013. 'Age Discrimination'. In *International Encyclopedia of Ethics*. American Cancer Society. <https://doi.org/10.1002/9781444367072.wbiee284>.
- Daniels, Norman. 1988. *Am I My Parents' Keeper?: An Essay on Justice between the Young and the Old*. Oxford University Press.
- Derous, Eva, and Jeroen Decoster. 2017. 'Implicit Age Cues in Resumes: Subtle Effects on Hiring Discrimination'. *Frontiers in Psychology* 8. <https://doi.org/10.3389/fpsyg.2017.01321>.
- Gosseries, Axel. 2014. 'What Makes Age Discrimination Special? A Philosophical Look at the ECJ Case Law'. *Netherlands Journal of Legal Philosophy* 43 (1): 59–80.
- Lippert-Rasmussen, Kasper. 2014. *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*. Oxford; New York: Oxford University Press.
- . 2017. 'The Philosophy of Discrimination: An Introduction'. In *The Routledge Handbook of the Ethics of Discrimination*, edited by Kasper Lippert-Rasmussen. London ; New York: Routledge.
- Macnicol, John. 2006. *Age Discrimination: An Historical and Contemporary Analysis*. <https://doi.org/10.1017/CBO9780511550560>.
- Malinen, Sanna, and Lucy Johnston. 2013. 'Workplace Ageism: Discovering Hidden Bias'. *Experimental Aging Research* 39 (4): 445–65. <https://doi.org/10.1080/0361073X.2013.808111>.
- McKerlie, Dennis. 1992. 'Equality Between Age-Groups'. BRL:KIE/38230. 1992. <https://repository.library.georgetown.edu/handle/10822/739945>.
- . 2001. 'Justice Between the Young and the Old'. *Philosophy & Public Affairs* 30 (2): 152–77. <https://doi.org/10.1111/j.1088-4963.2001.00152.x>.
- Moreau, Sophia. 2017. 'Discrimination and Freedom'. In *The Routledge Handbook of the Ethics of Discrimination*, edited by Kasper Lippert-Rasmussen. London ; New York: Routledge.
- Nussbaum, Martha C., and Saul Levmore. 2017. *Aging Thoughtfully: Conversations about Retirement, Romance, Wrinkles, and Regret*. Oxford University Press.
- Overall, Christine. 2006. 'Old Age and Ageism, Impairment and Ableism: Exploring the Conceptual and Material Connections'. *NWSA Journal* 18 (1): 126–37.
- Schauer, Frederick. 2017. 'Statistical (And Non-Statistical) Discrimination'. *The Routledge Handbook of the Ethics of Discrimination*. 23 August 2017. <https://doi.org/10.4324/9781315681634-5>.
- Shin, Patrick S. 2009. 'THE SUBSTANTIVE PRINCIPLE OF EQUAL TREATMENT'. *Legal Theory* 15 (2): 149–72. <https://doi.org/10.1017/S1352325209090090>.
- Wedeking, Gary A. 1990. 'Is Mandatory Retirement Unfair Age Discrimination?' *Canadian Journal of Philosophy* 20 (3): 321–34. <https://doi.org/10.1080/00455091.1990.10716494>.