

Age Discrimination: Is It Special? Is It Wrong?¹

Katharina Berndt Rasmussen

Stockholm University & Institute for Futures Studies

katharina.berndt@philosophy.su.se

PREPRINT – please refer to and cite only the published version

Abstract

In this chapter, I examine the moral status of age discrimination by bringing together accounts of the wrongness of discrimination with accounts of the specialness of age (as opposed to, e.g., gender or race). I provide an overview of the main specialness considerations concerning age-based treatment and suggest a taxonomy for different forms of age discrimination. I then relate three influential accounts of the prima facie wrongness of discrimination to this taxonomy. I show how these accounts differ in their moral assessment of different forms of age discrimination. I moreover explore different ways in which specialness considerations can affect the prima facie and overall moral status of age-based treatment, on these different accounts.

1 Introduction

Imagine Dana and Eli, two applicants for a vacant position. While both are qualified in all relevant respects, neither is called to be interviewed for the job—simply because of their gender (Dana is female) and race (Eli is black), respectively. These cases are intuitively clear instances of discrimination, as well as intuitively morally wrong.

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

An agent (say, an employer) group discriminates against someone (say, an applicant) on grounds of property P, by doing something (e.g., disregarding their application) if and only if:

- (i) in disregarding the application, the employer treats the applicant *worse* than she would have treated him, had he not had P, or had she not believed him to have P,
- (ii) it is *because* the applicant has P or because the employer believes that he has P, that she treats him worse, and
- (iii) 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.²

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

¹ Many thanks to the participants of the Prague (2017) and Madrid (2019) T.A.D.S. workshops and to an anonymous reviewer for valuable comments. Special thanks to A. Gosseries and G. Bognar for editorial support, constructive comments, and never-ending patience. Financial support by The Swedish Research Council, Riksbankens Jubileumsfond, and Marianne & Marcus Wallenberg Foundation is gratefully acknowledged.

² Cf. (Berndt Rasmussen 2019), (Lippert-Rasmussen 2014) and many of the entries in (Lippert-Rasmussen 2017). In *legal* terms, this definition captures (possibly legally acceptable) differential treatment as well as (unlawful) discrimination. On the *moral* status of these phenomena, see section 4.

- (a) being made worse off (e.g., deprived of an opportunity) or
- (b) being treated as inferior (e.g., considered not worthy of equal consideration with other—male or white—applicants).³

And arguably, treating someone worse in at least one of these senses, just because of their socially salient property, is *prima facie* morally wrong.⁴ Hence, *ceteris paribus*, Dana and Eli are wrongfully discriminated against.

Now imagine Alex, Billie, and Charlie, three equally qualified applicants for a vacant position. They belong to three different birth cohorts (say, three consecutive generations) and thus to different age groups: Alex is 60+, Billie is 40, and Charlie is under 18. While Billie is called back 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 many societies, and in this sense comparable to being female and being black. Moreover, in not being called back, Alex and Charlie suffer the same form of worse treatment as Dana and Eli. Hence, even their cases should be classified as discrimination by the above definition and count as morally wrong for the same reasons.

Writers on age discrimination, however, suggest that age is ‘special’, i.e., relevantly different from other grounds of discrimination, such as gender or race. Such specialness in turn might have moral ramifications, possibly making age discrimination less severe, or more justifiable, than these other forms. The specialness of age may thus translate into lesser moral seriousness of age discrimination. Our 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, there might be good reasons to change these practices and discount these intuitions.

This chapter examines the moral status of age discrimination by bringing together accounts of the wrongness of discrimination with accounts of the specialness of age. Section 2 summarises the special features of age and their role in different proposals to justify the use of age criteria and suggests a template within which these specialness considerations become relevant: the argument from specialness. Section 3 explores different forms of age discrimination. Section 4 presents three influential accounts of the wrongness of discrimination and shows that different forms of age discrimination are covered by different wrongness accounts. I moreover return to the proposals justifying the use of age criteria, based on the specialness of age, and explore the roles they can play under these different accounts. Section 5 concludes.

³ See (Berndt Rasmussen 2019).

⁴ For ease of exposition, I switch between talking about an action’s being *prima facie* (morally) wrong and there being a *prima facie* (moral) reason against the action.

2 The specialness of age and age-based treatment

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 (henceforth without the qualifier) has three special features, as compared to, e.g., gender or race:

- (A) perfect *passage-of-time correlation* (by definition);
- (B) *high divisibility*: a normal life span can be partitioned into many alternative age segments of different sizes; and
- (C) *multiple belonging*: a person with a normal life span will belong to a number of such age segments as time passes.⁵

A key idea is that, because of these special features, relying on age criteria (e.g., age limits) may sometimes be justified in two different ways. Firstly, means–end efficiency: age is a good proxy (statistical indicator) for certain given target variables and as such a tool for more efficient decision making. Secondly, overall betterness: using age criteria leads to better overall outcomes, in some sense. I will discuss these two proposals in the following two subsections, indicating which specialness features (A–C) become relevant at different points in the arguments.

2.1 Means–end efficiency: age as *reliable* and *precise proxy*

In typical employment cases, applicant age is 'significantly correlated' (Gosseries 2014: 63f.) to, e.g., job qualifications and expected productivity. This is mainly due to specialness feature (A), passage-of-time correlation.

To illustrate: Charlie, who is under 18, 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. Also, under-18-year-olds are likely not to have the cognitive capacities required for jobs where risk management is central. This has to do with typical neurological/brain development. Alex, who is 60+, is likely to lack some of the physical and cognitive abilities of an applicant in their 40's, like 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.

In a medical context, given statistical facts about life expectancy, age can work as a proxy for remaining life years (Bognar 2015; Cupit 2013). Thus, age can be one relevant factor for allocating scarce health resources, as exemplified by the concepts of quality-adjusted life year (QALY), and disability-adjusted life year (DALY).⁶

Thus, for a variety of reasons, age can be 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 an individual assessment. Yet, since the latter tend to be more time-consuming, complicated, or morally

⁵ For discussions of these features, see (Bidadanure 2016); (Gosseries 2014); (Cupit 2013); (Macnicol 2006).

⁶ See (Bickenbach 2016).

problematic, relying on age criteria can facilitate decision-making. Thus, to promote means–end efficiency, there may be reasons to use age criteria.

Justifying age criteria by reference to means–end efficiency also relies on specialness feature (B), high divisibility. Since a normal lifespan is divisible into any set of age segments that are relevant in a given context, using age criteria potentially allows for higher levels of ‘precision’ (Gosseries 2014: 63f.; cf. Bidadanure 2016) in approximating the relevant target variables, compared to, e.g., gender and race. Using more *precise* age criteria allows for increased means–end efficiency (at least if information is not too costly).

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

Whether something is a good 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., 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. This means that, if we seek to justify the use of age criteria by appealing to the fact that age is a good 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 (Gosseries 2014: 65f.; Lippert-Rasmussen 2014: 283–299). The analogy with gender and race thus serves to illuminate the point that establishing age as a reliable and precise proxy for certain target variables is not *sufficient* for an overall justification of age criteria.

We should, moreover, pause to note that this indeed is not *necessary* either. One could reject the idea that age is a reliable, sufficiently precise proxy for specific individual target variables, and still concede that relying on age criteria allows for better overall outcomes on a collective level. 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 may be used primarily as means for achieving a better collective outcome, even though correlation with specific *individual* target properties may be weak. The concern for better overall outcomes is thus distinct from the concern for means–end efficiency. And overall outcomes need to be considered even in the final analysis of the justificatory force of means–end efficiency considerations.

2.2 Better overall outcomes: *utility* or *fairness/equality*

Attempts to justify age criteria from better overall outcomes are mainly concerned either with utility or with fairness/equality. One key consideration under the headline of *utility* is collective ‘sequence efficiency’ (Gosseries 2014: 70f.): organising certain aspects of society around age criteria promotes total utility, as it makes society overall more functional. The education system is an obvious example. There are efficiency reasons for making people go

through stages of education, starting at a young, learning-conducive age and consecutively taking higher levels, prior to entering the labour market.

Again, however, one could object that there might be overall efficiency gains from race or gender discrimination—say, in a caste-like hierarchy—which we would not grant any justificatory force. With age, though, there is a crucial difference. In the education system example, *collective* sequence efficiency originates from *individual* sequence efficiency.⁷ In certain contexts, ordering someone's activities in chronological sequences promotes her achievement of the context-given objectives, due to specialness feature (A): 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 specialness consideration—and one which can appeal even to non-utilitarians. The trade-offs it is based on are not interpersonal, as in the sexist or racist caste system, where some people's extra burdens generate benefits for others. Rather, the trade-offs are intrapersonal: collective gains result from a more efficient distribution of burdens and benefits over individuals' lifetimes, as allowed by specialness feature (C), multiple belonging. Moreover, specialness feature (B), high divisibility, facilitates the tailoring of age segments to improve intrapersonal distributions—to the overall benefit of both the individuals and the collective whole.

From a perspective of *fairness or equality*, age criteria seem problematic as they imply worse treatment (in at least one of the above senses) of some—and hence inequality. Old Alex and young Charlie are denied an employment opportunity, which is granted to middle-aged Billie. The key idea with justifying such treatment, nevertheless, from a fairness or equality perspective is that, in the case of age, the inequality is mitigated once we consider whole lives.

According to 'complete life neutrality' considerations (Gosseries 2014: 66ff), for a fixed set of age criteria, *ceteris paribus*, there will be no inequality in treatment over complete lives. The *ceteris paribus* clause ensures that—due to specialness feature (C), i.e., multiple belonging—each age criterion for resource allocation impacts equally on everyone, if only at different points in time, thus treating each equally over their lifetime.⁸

However, this idea presupposes a very demanding *ceteris paribus* condition, where everyone has the same life length as well as the same lifetime profile of needs, goals and desires. Yet realistically, 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 65 benefits only whoever lives that long (Cupit 2013). Moreover, changes in these age

⁷ Bidadanure (2016: 247) calls this 'lifespan efficiency'. Alternatively one could ask which resource allocation pattern over her lifetime a rational chooser would prefer from behind a Rawlsian veil of ignorance (Cupit 2013); cf. (Bognar 2015); (Bognar 2008); (Daniels 1988).

⁸ Some egalitarians may object that *relational* inequalities at any specific time matter and do not disappear because of reversed inequalities at some other time (Bidadanure 2016).

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.⁹

According to ‘affirmative egalitarian’ considerations (Gosseries 2014: 70ff), age criteria may effectively *reduce* existing 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 jobs are made vacant and filled with members of the subsequent cohorts, shifting employment benefits to some previously unemployed. Later on, these cohorts of course have to do the same for their successors. Then, people’s life trajectories are equalised in the sense that fewer will go through stretches of unemployment *prior to*, while all are retired *after*, this age threshold.¹⁰

Another example is the allocation of life-saving treatment: fairness may require 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 someone beyond this threshold (Bognar 2008). Similarly, 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.

Acknowledging inequalities over whole lives—e.g., in employment opportunities, or life years—combined with special feature (C), multiple belonging, may thus provide affirmative egalitarian support for age-based treatment.

2.3 The argument from specialness

To pinpoint the role of specialness considerations for the overall moral assessment of age discrimination, I propose the following *argument from specialness*:

- (1) There is a prima facie reason *against* group discrimination.
- (2) Age-based treatment is a form of group discrimination (like gender- or race-based treatment).
- (3) In some contexts, specialness considerations concerning means–end efficiency or overall better outcomes provide reasons *for* age-based treatment that outweigh the prima facie reason against it.
- (4) Hence, in such contexts, we have overall reason *for* age-based treatment.

The specialness considerations of premise (3) have just been explored. In section 3, I briefly address premise (2) and consider four ways in which age-based treatment constitutes group

⁹ The ‘diversification’ approach (Gosseries 2014: 63) highlights that complete life inequalities, resulting from age criteria, are *moderated* due to age’s specialness: criteria tend to vary across contexts (due to (B)), and resulting disadvantages tend to spread out over an individual’s lifetime (due to (C)).

¹⁰ See (Arneson 2006: 797f.); (Wedeking 1990). For dissenting views, see (Overall 2006); (Nussbaum and Levmore 2017).

discrimination. In section 4, I consider three influential accounts of the prima facie wrongness of group discrimination, concerning premise (1).

3. Forms of age discrimination

Whenever an agent's age-based treatment of another individual relies on a statistical correlation between age and some agent-relevant end—aiming at means-end efficiency—we are dealing with *statistical discrimination*. We can capture the statistical version of age discrimination by modifying condition (ii) of our definition of group discrimination:

(ii*) it is because the employer believes that the applicant has property P, *and that P is statistically correlated to a target variable that is relevant for the employer*, that she treats him worse.¹¹

The general idea is that an agent believes that some socially salient property (say, old age) is statistically correlated to a relevant end (such as productivity) and therefore treats old applicants worse in the sense of making them worse off or treating them as inferior (e.g., by straightforwardly dismissing their application). Note that agents and subjects could also be collectives, rather than individuals: say, a society enacting policies that affect age groups differently, on the basis of widely shared or publicly communicated beliefs concerning such statistical correlations.

Statistical discrimination is typically considered as an instance of direct discrimination (Lippert-Rasmussen 2014: 88). Direct discrimination is often characterised in terms of an agent's *intention to treat people differently*—and indirect discrimination in terms of a 'facially neutral' policy that treats all the same, but *non-intentionally* has *disparate impact*.¹²

This orthodox dichotomy, however, confounds *two* distinctions. One is between differential treatment (treating older and younger people differently) and disparate impact (treating all the same, when this affects older and younger people differently). The other distinction is between the intentional and non-intentional use of, e.g., age-related criteria. These two distinctions result in four, rather than two, distinct forms of discrimination (Berndt Rasmussen 2020):

	<i>Differential treatment</i>	<i>Disparate impact</i> of the same treatment
<i>Intentional</i>	(I) Intentionally using age as a criterion for differential treatment: setting a minimum hiring age of 25.	(II) Employing a facially neutral criterion, 'overqualified', while intending its differential impact on older applicants.
<i>Non-intentional</i>	(III) Treating differently without intention: perceiving equally qualified older applicants as less qualified, due to implicit age bias. ¹³	(IV) Employing a facially neutral criterion, 'overqualified', without intending its differential impact on older applicants.

Table 1: four forms of discrimination with age-related examples

¹¹ See (Schauer 2017: 42f.); (Lippert-Rasmussen 2014: 81).

¹² See (Lippert-Rasmussen 2017, 3); (Moreau 2017, 166–67); (Altman 2016, §3.1); (Macnicol 2006, 21).

¹³ For empirical studies, see (Derous and Decoster 2017); (Malinen and Johnston 2013); (Cortina et al. 2013).

Age discrimination, understood as statistical discrimination according to condition (ii*), can fall under either differential treatment or disparate impact on the intentional dimension. E.g. (Table 1, box I), an employer who sets a minimum hiring age of 25 due to neurological evidence that under-25-year-olds tend to have underdeveloped brain regions for risk assessment, statistically discriminates against those under 25 in the differential treatment sense. (Had the employer set this age limit because of her dislike of younger people, this would not be *statistical* discrimination.) Or (box II), an employer who uses the criterion ‘overqualified’ because she prefers younger employees can be said to statistically discriminate against older applicants, in the disparate impact sense, if this preference is due to her belief that younger age correlates with ‘more grit’. (Had her preference instead reflected a fear of aging, this would not be *statistical* discrimination.)

Even collective agents can be said to statistically discriminate (against individuals or groups), e.g., when a society publicly justifies a minimum employment age with reference to statistical correlations between such a policy and overall education levels. (Had this policy instead been publicly justified by reference to religious doctrine, this would not be *statistical* discrimination.)

According to this analysis, all statistical age discrimination is intentional (in an individual or collective sense). But not all forms of intentional age discrimination fall under the label of statistical discrimination: e.g., those due to dislike, fear, or religious doctrine. And, as *Table 1* suggests, there are also forms of age discrimination that constitute non-intentional discrimination: e.g., due to implicit age bias (box III), or due to failing to consider that one’s criteria may affect different age groups differently (box IV).

A further upshot is that, by teasing apart the two conflated distinctions underlying the orthodox dichotomy of direct and indirect discrimination, we get equipped with a wider, yet still unifying lens. The lens is wider, in helping us capture a wider variety of phenomena that qualify as age discrimination. It is unifying, in providing a coherent template for the further evaluation of this variety. Using this new lens means that we refrain from using the conflated direct/indirect distinction, and instead try to state clearly which of the four form of discrimination we are interested in or dealing with in a specific case.

4. The wrongness of age discrimination

This new taxonomy has repercussions for the moral status of discrimination. I will sketch three prominent accounts of the prima facie wrongness of discrimination and apply them to the four forms of age discrimination. For each of these three accounts, I moreover consider whether reasons derived from age’s specialness, as identified in section 2, may outweigh such prima facie reasons.

4.1 Intention-focused accounts

According to one influential account, an instance of discrimination is prima facie morally 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 (Alexander 1992, 161). According to another such account, such an instance is prima facie wrong when and because it is ‘done from unwarranted animus or prejudice against persons of [a socially salient] type’

(Arneson 2006, 787f.). These accounts thus focus on the agent's *intentions*: the mental states motivating her act.

To illustrate: if 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 to an aversion against older people (Table 1, box I), this constitutes *prima facie* wrongful age discrimination. The same holds if Alex' application instead fails some facially neutral criterion (say, 'overqualified') that the employer has posited because of such mental states (box II). Now, consider some variations. E.g., the employer may be motivated by affirmative egalitarian considerations: she wants to give an equalising push to younger applicants, who are not as firmly established in the job market as their older competitors. Or she goes for means-end efficiency and uses age as a proxy, thus engaging in statistical discrimination. Then, the employer's treatment of Alex turns out as *not* *prima facie* wrong, since it is motivated not by the above objectionable intentions, but rather by specialness considerations. As sketched so far, intention-focused accounts provide intuitively correct judgments of a variety of cases, and moreover can take specialness considerations into account.

However, now suppose that the employer acts—*intentionally*—from religious doctrine that does not imply any assumptions of inferiority, aversion, or prejudice. E.g., she believes that people above 60 should be given time to rest or worship, instead of having to work (Table 1, box I). Turning down Alex's application would then *not* qualify as *prima facie* wrong.¹⁴ Moreover, all *non-intentional* forms of age-based treatment—due to implicit age bias (box III) or mere failure to grasp the hiring criteria's disparate impact (box IV)—get a moral free pass. Since they are not motivated by the objectionable intentions, according to these accounts there is no *prima facie* reason against them. These verdicts seem intuitively less plausible. Proponents of intention-focused accounts must either bite these bullets—or considerably revise their accounts to include even 'benign' intentions and non-intentional forms of discrimination (e.g., framing implicit biases as aversions; or deriving the failure to consider the elderlies' vulnerabilities from beliefs about their inferiority).

Now, consider only cases which *do* qualify as *prima facie* wrongful age discrimination, according to intention-focused accounts. Proponents have the further option of accepting the argument from specialness, such that the *prima facie* reasons against the treatment in question can be outweighed by other reasons, due to age's specialness. E.g., while the employer's decision not to call Alex back because of ageist prejudice is *prima facie* wrong, affirmative egalitarian considerations may provide an outweighing reason in favour of the decision, all things considered. However, proponents of intention-focused accounts will presumably object that bad intentions cannot be counterbalanced by good outcomes. They then deny premise (3), thus restricting specialness considerations to the *prima facie* level.

Intention-focused accounts thus manage to capture our intuitions in some cases of intentional age-based treatment but fail in others. They moreover cannot differentiate between non-intentional cases, but give them all a moral free pass. Finally, while they can account for specialness considerations when these constitute exculpating mental content on

¹⁴ See (Arneson 2006: 801f.)

part of the agent, they will presumably deny them the weightier status of outweighing moral reasons.

4.2 Social meaning accounts

While the previous accounts focus on the intentions behind the act, the present ones focus on its meaning for others. Here, 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 (Hellman 2008: 36).¹⁵ To assess whether these conditions are met, we need to interpret the act within its social context by examining conventional methods of expressing respect and disrespect, the agent’s social status, etc. Such an assessment will take certain mental states of people—which may include the agent, but also the victim, the general public, possibly a hypothetical impartial observer—into account: beliefs and attitudes concerning, e.g., socially salient groups, power relations, and social norms.¹⁶ But in contrast to the previous accounts, the mental states of the agent are not alone decisive here.

Social meaning accounts have the potential to assess all four forms of age discrimination. E.g., regardless of whether the employer’s decision to disregard Alex’ application is motivated by blatant ageist prejudice or (Table 1, boxes (I) or (II)) or due to unnoticeable implicit bias (box III), it is *prima facie* wrong if it expresses disrespect. Or, if the employer uses the facially neutral criterion ‘overqualified’, being genuinely unaware of how this affects older applicants (box IV), her action may still express disrespect if it signals that older people’s interests may safely be neglected. Note that even seemingly benign actions may be rendered objectionable if they become tainted by disrespect. E.g., if a volunteer helps the elderly with an air of superiority or a patronising demeanour, this may then render her treatment demeaning, via the uptake of others, and hence *prima facie* wrong. In the abovementioned cases, social meaning accounts deliver intuitively plausible verdicts.

On the other hand, some intuitively objectionable age-based treatment might be exculpated. Consider an employer who immediately dismisses an older, clearly qualified applicant as unqualified. In a social context where everyone is aware of the ubiquity of implicit age bias, the action may then be exculpated as a merely unfortunate consequence of the employer’s all-too-human psychology rather than an expression of disrespect. This seems more problematic.

The idea that the moral status of age-based treatment hinges on its social meaning within a given context has further troubling implications. Imagine a social context with commonly shared ageist beliefs, where age-based treatment is routinely seen as a ‘natural’ consequence of age-based properties. No such treatment may then even register as *prima facie* wrong, since it never expresses disrespect. In the extreme, in a fully age-segregated society there may be no instance of wrongful age discrimination at all. This seems intuitively questionable, to put it mildly (consider equivalent examples involving gender or age).

¹⁵ See (Bidadanure 2016); (Shin 2009).

¹⁶ According to Hellman (2008: 75), demeaningness is *modestly* objective: neither entirely independent of people’s beliefs and practices (strongly objective), nor totally dependent on what the majority thinks (minimally objective).

Now, consider only cases which *do* qualify as prima facie wrongful age discrimination according to social meaning accounts. Proponents may, again, accept or reject that prima facie reasons against age discrimination can be outweighed by specialness considerations, as proposed in premise (3) of the argument from specialness. Here, it seems rather plausible that they should accept the argument. Specifically, premise (3) captures the dynamics involved in changes to the social meaning—and thus arguably the moral status—of age-based treatment. Consider the employer’s decision not to call back Charlie, who is under 18 but clearly qualified. In a social context where this expresses lack of respect for youngsters this is then prima facie wrongful age discrimination. Now assume that public debate arises around the social benefits of a minimum employment age of 18. Advocates point out that this would ensure that youngsters finish their basic education prior to entering the labour market, thus referring to collective sequence efficiency considerations. This means that specialness considerations are employed as outweighing reasons in favour of an age limit, along the lines of premise (3). If this debate eventually yields widespread agreement on this age limit, the prima facie status of the type of act changes: the employer’s next decision not to call back a qualified under-18-year-old applicant, just because of their age, will adhere to changed social norms and thus no longer express disrespect.

This reasoning along the lines of the argument from specialness brings out that social meaning accounts have the advantage of being responsive to changing social contexts, thus keeping their verdicts closer to common intuitions (within these contexts). The flipside is that these accounts risk losing the critical distance to the social context they are supposed to evaluate—as seen in their counterintuitive assessments within the above extremely age-segregated society.

4.3 Harm accounts

Both previous accounts of the wrongness of discrimination pick up on one sense of ‘being treated worse’: (b) being treated as inferior. The present account picks up on the other: (a) being made worse off. Here, an instance of discrimination is prima facie wrong when and because it harms the victim, i.e., *makes them fare worse* as compared to what they would have, had they not had socially salient property P (Berndt Rasmussen 2019; Lippert-Rasmussen 2014). Thus, harm accounts compare the actual and counterfactual outcome, relative property P, in terms of the subject’s welfare levels.

Harm accounts also have the potential to assess all four forms of discrimination. Consider Alex who is 60+ and not called back. Harm accounts are insensitive to whether the decision is motivated by open ageism (Table 1, box I) or implicit bias (box III), whether a facially neutral criterion is used in a premediated way (box II) or due to lack of better understanding (box IV). What counts is only the resulting harm to Alex, that she (in the relevant sense) fares worse. But does she, really?

As the evidence from middle-aged Billie indicates, had Alex been middle-aged, she would have been called back. Has she thus 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.

Arguably, had Alex been middle-aged at this point in time (like Billie), she would have been 60+ at some later point—and, *ceteris paribus*, 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 sequence efficiency give us reason to doubt that such a claim would eventually bear out. Thus, under strict *ceteris paribus* conditions with individually (complete life) beneficial age criteria, the employer's decision not to call Alex back due to her age does not make Alex worse off and is thus not *prima facie* morally wrong.

In many cases, however, such demanding *ceteris paribus* conditions do not hold (e.g., employment opportunities vary over birth cohorts), or age criteria do not achieve individual sequence efficiency. Then, the employer's decision may make Alex worse off and thus be *prima facie* morally wrong.

Thus, for harm accounts, *some* specialness considerations—individual sequence efficiency, complete life neutrality—are relevant on the *prima facie* level. Others are not: e.g., the affirmative egalitarian claim that that mandatory retirement rules reduce inequalities between people will not matter for the assessment of whether Alex is harmed. The same goes for collective sequence considerations (that focus on collective benefits) or means-end considerations (that focus on the employer's benefit of using age as a proxy). Thus, while all specialness considerations can be relevant at the *prima facie* level for intention-focused accounts (if they figure in the agent's mental content) and for social meaning accounts (if they shape the meaning of the act), harm accounts are more selective at this level.

However, even here, all specialness considerations may re-enter the picture in the role specified by premise (3) of the specialness argument, generating outweighing reasons for age-based treatment that is *prima facie* wrong. E.g., the employer's decision not to call Alex back may be harmful to Alex and hence *prima facie* wrong, but nevertheless overall justified due to the promotion of utility (through collective sequence efficiency) or equality (through affirmative egalitarian effects).

One final note: the central focus of harm accounts is the *victim* and the discriminating act's effects on how she fares in life, which is intuitively appealing. Some might worry that harm accounts nevertheless miss other important factors, as covered by the two rival accounts, respectively. Still, depending on how the concept of harm is spelled out, (part of) the victim's harm may result from, e.g., her perception of the action as socially demeaning or of the agent's intentions. The harm account may thus be able to integrate some of these factors as well.

5 Conclusion

In this chapter, I have brought together considerations of the specialness of age with three influential accounts of the moral wrongness of discrimination. I have evaluated these accounts' moral verdicts for age-based treatments, and explored how they can take

specialness considerations into account. In doing so, I have proposed a template for recognising, understanding, and morally evaluating different forms of age discrimination.

As a final concession, the chapter does not give determinate answers to its two title questions. My hope is that this proposed template may be useful for others interested in further pursuing these questions.

References

- Alexander, Larry. 1992. 'What Makes Wrongful Discrimination Wrong Biases, Preferences, Stereotypes, and Proxies'. *University of Pennsylvania Law Review* 141 (1).
- 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).
- Berndt Rasmussen, Katharina. 2020. 'Implicit Bias and Discrimination'. *Theoria*, 86 (6).
- Berndt Rasmussen, Katharina. 2019. 'Harm and Discrimination'. *Ethical Theory and Moral Practice* 22 (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).
- Bognar, Greg. 2008. 'Age-Weighting'. *Economics & Philosophy* 24 (2).
- Bognar, Greg. 2015. 'Fair Innings'. *Bioethics* 29 (4).
- 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).
- Cupit, Geoffrey. 2013. 'Age Discrimination'. In *International Encyclopedia of Ethics*, edited by Hugh LaFollette. Blackwell Publishing.
- Daniels, Norman. 1988. *Am I My Parents' Keeper?: An Essay on Justice between the Young and the Old*. Oxford: Oxford University Press.
- Derous, Eva, and Jeroen Decoster. 2017. 'Implicit Age Cues in Resumes: Subtle Effects on Hiring Discrimination'. *Frontiers in Psychology* 8.
- Gosseries, Axel. 2014. 'What Makes Age Discrimination Special? A Philosophical Look at the ECJ Case Law'. *Netherlands Journal of Legal Philosophy* 43 (1).
- Hellman, Deborah. 2011. *When is Discrimination Wrong?* Cambridge, MA: Harvard University Press.
- Lippert-Rasmussen, Kasper. 2014. *Born Free and Equal?: A Philosophical Inquiry into the Nature of Discrimination*. Oxford: Oxford University Press.
- Lippert-Rasmussen, Kasper. 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*. Cambridge: Cambridge University Press.
- Malinen, Sanna, and Lucy Johnston. 2013. 'Workplace Ageism: Discovering Hidden Bias'. *Experimental Aging Research* 39 (4).

- McKerlie, Dennis. 1992. 'Equality Between Age-Groups'. *Philosophy & Public Affairs* 21 (3).
- 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: Oxford University Press.
- Overall, Christine. 2006. 'Old Age and Ageism, Impairment and Ableism: Exploring the Conceptual and Material Connections'. *National Women's Studies Association Journal* 18 (1).
- Schauer, Frederick. 2017. 'Statistical (And Non-Statistical) Discrimination'. *The Routledge Handbook of the Ethics of Discrimination*, edited by Kasper Lippert-Rasmussen. London; New York: Routledge.
- Shin, Patrick S. 2009. 'The Substantive Principle of Equal Treatment'. *Legal Theory* 15 (2).
- Wedeking, Gary A. 1990. 'Is Mandatory Retirement Unfair Age Discrimination?' *Canadian Journal of Philosophy* 20 (3).