

AFFIRMATIVE
ACTION



A VIEW FROM
THE GLOBAL SOUTH

EDITORS

OCKERT DUPPER
KAMALA SANKARAN


SUN PRESS

Affirmative Action: A view from the global South

Published by SUN MeDIA Stellenbosch under the imprint SUN PRESS.

All rights reserved.

Copyright © 2014 AFRICAN SUN MeDIA and the editors/authors

No part of this book may be reproduced or transmitted in any form or by any electronic, photographic or mechanical means, including photocopying and recording on record, tape or laser disk, on microfilm, via the Internet, by e-mail, or by any other information storage and retrieval system, without prior written permission by the publisher.

First edition 2014

ISBN 978-1-920689-46-9

ISBN 978-1-920689-47-6 (e-book)

Set in Arno Pro 12/14

Cover design: SUN MeDIA Stellenbosch

Typesetting: SUN MeDIA Stellenbosch

Cover photograph:

SUN PRESS is an imprint of AFRICAN SUN MeDIA. Academic, professional and reference works are published under this imprint in print and electronic format. This publication may be ordered directly from www.sun-e-shop.co.za.

Produced by SUN MeDIA Stellenbosch.

www.africansunmedia.co.za

africansunmedia.snaplify.com (e-books)

www.sun-e-shop.co.za

CONTENTS

PREFACE	1
---------------	---

PART I

COMPARATIVE, THEORETICAL AND INTERNATIONAL PERSPECTIVES

1. AFFIRMATIVE ACTION IN COMPARATIVE PERSPECTIVE	7
<i>Ockert Dupper</i>	
1.1 Introduction	7
1.2 Affirmative action and social identity	14
1.3 Demarcation of beneficiaries	21
1.4 Impact assessment	27
1.5 Justification and termination	34
1.6 Concluding remarks	40
2. DEMOCRATIC PARTICIPATION IN MAKING AND ENFORCING AFFIRMATIVE ACTION SCHEMES	43
<i>Bob Hepple</i>	
2.1 Introduction	43
2.2 Reflexive regulation	45
2.3 The failure of command and control	50
2.4 The nature and extent of participation and accountability	58
2.5 Conclusion	60
3. THE CASE AGAINST THE CASE AGAINST AFFIRMATIVE ACTION	63
<i>Ana Alfinito Vieira & Alex Graser</i>	
3.1 Introduction	63
3.2 The broader context: AAPs and other EOPs	64

3

THE CASE AGAINST THE CASE AGAINST AFFIRMATIVE ACTION

Ana Carolina Alfinito Vieira & Alex Graser¹

3.1 Introduction

Affirmative action policies (AAPs) have spurred heated debates ever since they came into existence. They have been attacked² and defended³ over and over again, and yet,

- 1 Alex Graser is Professor of Public Law and Policy at the University of Regensburg and Fellow of the Hertie School of Governance, Berlin; Ana Carolina Alfinito Vieira is a doctoral fellow at the Max Planck Institute for the Study of Societies, Cologne. The present research was supported by grants from the Deutsche Forschungsgemeinschaft and MaxnetAging, a project of the Max Planck Society.
- 2 The 'case against affirmative action' has been made far more often than could be cited here. Suffice it therefore at this point that we pick two articles that were published under that very title and hence may serve as a point of reference for ours: Louis P Pojman 'The Case Against Affirmative Action', *International Journal of Applied Philosophy* 12, no. 1 (1998): 97–115; Terry Eastland, 'The Case Against Affirmative Action', *William and Mary Law Review* 34, no. 1 (1992): 33.
- 3 As the defences also abound, we limit ourselves to the very succinct attempt at rebutting the common objections by Stanley Fish, 'The Nifty Nine Arguments Against Affirmative Action in Higher Education', *The Journal of Blacks in Higher Education* 27 (2000): 79–81.

to this point, the issue does not seem to be settled. Indeed, it is quite unlikely that it ever will be, as it involves questions of *distributional* justice upon which political preferences might in most societies be irreconcilably divided. Nonetheless, there is still hope for at least an increased consensus on the desirability of this kind of policy – and thus a point in continuing the debate.

The present article sets out to advocate a more widespread use of AAPs by countering what we perceive to be a critical (set of) objection(s) to these policies. Namely, we refer to the argument that AAPs are unfair because they entail undue burden upon two classes of individuals that are affected by these policies: disadvantaged competitors and burdened addressees. We call this objection ‘critical’ because we think it is difficult to counter and implicit in most debates about AAPs. At the same time, it seems to be rarely addressed, let alone discussed thoroughly, neither by policy makers nor in the literature. It is our hope that by directly and systematically addressing this objection, we may help not only to clarify the terms of the debate surrounding AAPs, but also to further develop and consolidate the arguments in favour of these policies.

The basic structure of our argument is simple and comprises two steps. As there are many objections against AAPs, we will first filter out and develop what we consider to be the critical objection (3.3). In a second step, we will build up a counter argument that systematically tackles all facets of the critical objection, thereby seeking to weaken the case against AAPs (3.4). We shall end our contribution with a brief concluding section (3.5). Since we are dealing with a question of policy choice, it will be necessary to consider at the outset what alternative policy options exist besides AAPs to tackle inequality. Our chapter will therefore begin with a brief and introductory account of the broader array of equality-oriented policies (EOPs).

3.2 The broader context: AAPs and other EOPs

For most governments today, the promotion of equality – or conversely: the reduction of inequalities – is one of their central goals. Accordingly, they typically pursue a broad set of what we have here labeled EOPs.⁴ Of course, equality can mean different things, and the perceived obstacles to its achievement also vary across time and across polities. Therefore, the composition of the basket of EOPs deployed by any specific government at any given time also varies substantially.⁵

4 For an in-depth presentation of the concept of equality-oriented policies, see Alexander Graser, ‘Equality-oriented policies: A new concept in public policy?’ (18 February 2013), available at <http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2220289>.

5 For a comparative overview of EOPs in various jurisdictions, see Alexander Graser, Ana Alfinito Vieira & David Jackson, eds., *Equality-Oriented Policies – The Concept*, forthcoming with Nomos in 2014.

Despite such variability, we would like to introduce a categorisation that distinguishes in broad terms between three 'layers' of EOPs.⁶ The metaphor suggests that the policies belonging to these 'layers' have emerged roughly in the corresponding historical sequence. But we speak of 'layers' rather than 'waves' or 'stages' because they do not typically replace, but rather add to one another, so that, in fact, all three layers of EOPs typically coexist nowadays, and each of them continues to evolve. Also, there is a lot of variation across systems with regard to the exact shape of these layers.

The policies of the three layers differ with regard to the cause of inequality that they seek to tackle. First-layer EOPs aim at guaranteeing general equality before the law. They entail the abolition of status differences and the concurrent operation of a cross-cutting equality principle that is typically situated at some elevated rank within the normative hierarchy. While EOPs of this layer are fundamental, the kind of equality they implement remains formal. They do not target differences in material wealth among the members of a polity.

This is where second-layer EOPs come into play. These are the (re)distributive policies that seek to promote material equality within the polity and typically take the form of social benefits, taxes, or institutional spending. To the formal status equality, second-layer EOPs add materially equalised starting positions. Nonetheless, these policies do not achieve full equality in material terms across all individuals in a polity. In fact, they are not meant to do so, as it is widely accepted that differences in individual achievement or choice of lifestyle should be reflected in some degree of material inequality.⁷ Some inequalities that persist despite the operation of second-layer EOPs may hence be viewed as intended.

But this does not explain all persisting inequalities because second-layer EOPs, too, fail to address important drivers of inequality and are, therefore, also insufficient. It is not only differences in legal status or material endowment that cause inequality; inequality is also generated and regenerated at the level of interactions among individuals day by day. In fields such as employment relations, medical care,

6 For a more thorough exposition of this categorisation, see Alexander Graser and Ana Alfinito Vieira, 'Taming the Biased Black Box? On the Potential Role of Behavioral Realism in Anti-Discrimination Policy', forthcoming in: *Oxford Journal of Legal Studies*, (1) 2015; an earlier version of the paper (1 October 2012) is available at <http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2209999>.

7 For an exposition of the causes of inequality considered as legitimate by different approaches of political philosophy, see Gerald A Cohen, *Why not Socialism?*, (Princeton and Oxford: Princeton University Press, 2009). The position presented here is in line with what Cohen (at 18 ff.) calls 'socialist equality of opportunity'. This is the most egalitarian among the positions he presents. Hence, our statement in the text applies a fortiori for the other positions that Cohen distinguishes.

education, rental agreements, and many others, individual decisions can have a major impact on other people's situation. It is hence critical to target any bias with regard to skin colour, gender, age, or other features that may inhabit such individual decision-taking. This is what third-layer EOPs do – some of them by addressing the existence of widespread popular bias in general (for example, through educational programmes and public awareness campaigns) and others by intervening on the individual level by targeting individual decisions and the processes through which they are made.

It is the latter type of third-layer EOPs with which this article is primarily concerned. They, in turn, comprise a variety of policies. Among others, these include, first, rules that prohibit certain criteria from being taken into account in the decision-making process – that is, anti-discrimination laws (ADLs). Second, they also include rules that prescribe a certain outcome to some or all decisions of a certain type. AAPs often fall in this second category.

It is apparent that these policies may interfere with individual autonomy. That is the price of intervening in individual decision-taking. At the same time, it should be noted that they do so with different intensities. ADLs would appear to be particularly respectful of individual autonomy as they only prohibit the consideration of specific criteria while otherwise leaving individual preferences untouched. This may be a reason that ADLs seem now to be the most widespread among third-layer EOPs. Across many systems, we have witnessed an impressive expansion of these policies over the last decades. This pertains to both the criteria they target as well as the realm of social life to which they apply.⁸

At this point, we conclude our brief survey of EOPs. We shall return to the observation of the predominance of ADLs again later on when we consider their relative attractiveness compared to AAPs. First, however, our focus in the ensuing sections will be on the policies that are central to this article.

3.3 What may be (particularly) wrong about AAPs

In this section, we will spell out what we perceive to be the core objection raised against AAPs – namely, their perceived unfairness. For this purpose, we shall sketch the different types of objection that can be raised against public policies in general and position the unfairness objection within this context (3.3.2). We will then further specify this objection and explain why it requires us to distinguish between

⁸ See Sandra Fredman, *Discrimination Law*, 2nd ed. (New York: Oxford University Press, 2011), for a thorough account of the historical development within the context of EU and UK law (38 ff.) and a systematic normative discussion of the proper scope of discrimination law (at 109 ff.).

two categories of individuals: disadvantaged competitors and burdened addressees (3.3.3). But as a first step, we shall clarify what exactly we mean by AAPs (3.3.1).

3.3.1 Defining AAPs

The term affirmative action is widely used in different regional and disciplinary contexts. It commonly denotes policies that are meant to help members of minority groups or other disadvantaged groups to get better access to various spheres of social life such as education, employment and political representation. The exact boundaries of the term may vary depending on the specific context. In the following discussion, we will give the term a somewhat narrower but also more precise meaning, and will understand AAPs as legal rules that target individual decisions and require (or in some circumstances just allow) decision-takers to afford preferential treatment to a certain group of individuals who are subject to these decisions. Prominent examples are quotas, say, for women in political or managerial leadership positions, or rules that afford members of disadvantaged ethnic groups a competitive advantage within formalised assessment schemes, be it for hiring purposes, for university access or tenders for public contracts.

3.3.2 The unfairness objection in context

In our perception, such AAPs face continued and at times fierce opposition in many jurisdictions. As indicated before, the critical objection against them seems to be that they are unfair. This may at first appear to be a rather unspecific statement. But in fact, it serves to exclude a broad range of potential objections from the scope of the present discussion.

This is because unfairness is only one of multiple reasons why a policy may be judged as bad.⁹ One might disagree with a policy's ends, or on the means that are employed to achieve them. With regard to the means, one may focus on the output and conclude that the policy is ineffective, causes unforeseen side-effects, or is even counter-productive. Or, looking also at input, it may appear that a policy is too costly or at least inefficient.

This is true not only for policies in general, but can also be applied to AAPs in particular. In principle, all of these objections could be launched at AAPs. Whether they obtain can hardly be discussed in the abstract as this will ultimately depend on the specific policy and its context. But to the extent that generalisation is possible, our impression is that, typically, neither the ends of AAPs nor their implementation is particularly contested.

⁹ Stuart Nagel, *Handbook of Public Policy Evaluation*, (Thousand Oaks: Sage Publications, 2011). See also Albert Hirschman, *Rhetoric of Reaction*, (Cambridge: Harvard University Press, 1991).

This is not to say, first, that everybody has always agreed on the need for AAPs. Any particular AAP may be considered unwarranted, be it because its beneficiaries do not really need that support, or because of a general hesitance about governmental intervention to achieve more equality. But it seems that AAPs are usually employed in situations in which there is a blatant under-representation of their beneficiaries. This makes consensus regarding the ends more likely.

Second, AAPs may, of course, also suffer from implementation deficits. It is hard to conceive of any kind of policy that could not, and AAPs are no exception. But it seems that AAPs can be framed in ways that render them hard to circumvent and hence relatively easy to enforce. The aforementioned examples testify to this: a fixed quota is a case in point, as is a minority bonus within a formalised application scheme. Implementation does not have to be a critical issue with these policies.

There is, to be sure, some debate about undesired side-effects of AAPs. The argument is that these policies may corroborate the negative public perception of their beneficiaries as a distinct group that is in need of special protection.¹⁰ While this is not implausible per se, it is highly doubtful that this potential long-term effect is strong enough to outweigh the short-term benefits to the group of beneficiaries. But this is mere speculation. Despite the advances that social psychology has made in the field of stereotype research,¹¹ there seems to be no prospect that this alleged side-effect could be verified in the foreseeable future.

In sum, there are, of course, multiple objections that can be raised against AAPs. But the ones reviewed so far do not seem particularly salient. They can hardly explain the level of controversy that has accompanied AAPs so far.

3.3.3 Specifying the unfairness objection

This takes us back to the initial proposition that the most important objection against AAPs is their perceived unfairness, which can now be further specified. We have said already that the problem is typically not that the beneficiaries of AAPs are not considered deserving of the support they are afforded. Hence, the perceived unfairness of AAPs primarily lies in the burden that AAPs entail for others.

However, the fact that there is a burden to others is not per se unusual. Most EOPs – and all of those that aim at some distributional effect – do entail such a burden. Whatever their beneficiaries receive must be taken from someone else. This is true

10 For a thorough discussion of this argument and its potential implications for policy making, see Ockert Dupper, 'Affirmative Action: Who, How and How Long?', *South African Journal on Human Rights* 24 (2008): 425; for a brief discussion see also Stanley Fish, 'The Nifty Nine Arguments Against Affirmative Action in Higher Education', 81.

11 On this research, see 3.4.2.2.

for any tax-financed social benefit to needy persons, and for all social insurance schemes that have a distributional dimension, and it is also for true income tax schemes, especially progressive ones.

3.3.3.1 *The perceived unfairness of individual burdens*

What is different in the case of AAPs is that there is a burden to others that is not spread across the entire polity within which the policy operates. In other words, there is a burden that is not a collective, but an individual one. AAPs may, to be sure, also bring about some collective burden. To the extent that leaving individual decisions untouched would lead to an efficient allocation of resources¹² – in this context, of human capital – the efficiency loss due to the intervention by way of an AAP could be viewed as a collective burden. But arguably, this is not the main burden that AAPs cause, and definitely, such collective burdens are nothing uncommon. What is special about AAPs is that they impose a significant burden upon some individuals.

3.3.3.2 *The two categories of unduly burdened individuals*

Among these individuals, we can distinguish two categories. One is that of burdened addressees, the other that of disadvantaged competitors. For purposes of illustration, imagine an employment context in which an affirmative action rule prescribes preferential hiring of a certain group. This will, at least in some cases, imply that the employer has to hire a candidate other than the one who is (judged to be) most qualified. The latter person we term the ‘disadvantaged competitor’, while the employer in this setting will be labeled ‘burdened addressee’. Obviously, the terms can be applied equally to other settings in which AAPs are used – such as, most notably, in admissions to tertiary education.

Both categories of persons are negatively affected, but in different ways. The disadvantaged competitor, first, is deprived of a benefit to which she is not entitled, to be sure, but which she would have enjoyed if the AAP were not in place. The above examples show how important this foregone benefit can be. Access to a certain job or a university may well have a significant effect on a person’s life.¹³

¹² On this argument, see 3.4.2.1.

¹³ For a pointed discussion of the situation of who we label as the ‘disadvantaged competitor’ and a critique of the individualistic premises underlying the presumption of disadvantage in that context, see Stanley Fish, ‘Reverse Racism, or How the Pot Got to Call the Kettle Black’, November 1993 (at subheading: ‘Why me?’ and below), available at <<http://www.theatlantic.com/magazine/archive/1993/11/reverse-racism-or-how-the-pot-got-to-call-the-kettle-black/304638/4/>>.

The burdened addressee, second, is imposed a choice that is not hers, and that – at least in her perspective – is also sub-optimal. The importance of this disadvantage is harder to gauge than that of the disadvantaged competitor. The infringement of the burdened addressee's autonomy may be viewed as more or less important from a principled perspective, but cannot be quantified. The loss related to the sub-optimal choice, by contrast, can better be assessed, albeit only for the specific case as it depends on the difference between the beneficiary of the AAP and the disadvantaged competitor. It may be marginal if they are almost equally qualified, but if not, it becomes more relevant.¹⁴ AAPs can and indeed often do address that latter issue by stipulating threshold criteria for the beneficiary to meet, thus ensuring that the difference is not too large.

3.4 Why this may not be so wrong after all

We have now spelled out what we consider the critical objection to AAPs. It is their unfairness against burdened individuals. When in the following discussion we turn to examining this objection, we shall do so from two perspectives that are distinct, but related.

The first is, what we will label the 'argument from individual equality' (3.4.1). This perspective is relevant for both categories of burdened individuals, because both suffer a disadvantage that is greater than that of all other persons in the polity in which the AAP operates. So, the problem here is that the AAP in effect singles out these burdened individuals for disparate negative treatment.

The other perspective is termed the 'argument from personal autonomy' (3.4.2). It is not relevant to the category of disadvantaged competitors, but only to that of burdened addressees. In the following, we will deal with both perspectives separately, but take account of their interdependence.

3.4.1 The argument from individual equality

We have said before that it is not unusual for policies that promote equality to do so in a fashion that distributes goods from some individuals to others; the benefit they afford to their beneficiaries comes at the expense of someone else. And we have also pointed out that one main problem with AAPs may be that this expense rests upon the shoulders of certain individuals rather than being borne collectively. This is the reason that AAPs, although intended to promote equality, might be viewed as resulting in an unequal treatment of the burdened individuals.

¹⁴ For a rule that allows for a wide differential of qualifications, see the South African Employment Equity Act 55 of 1998, especially the definition of 'suitable qualification' in s 20(3)(d), which stipulates that it may be sufficient if the beneficiary is capable of acquiring the abilities relevant for the job within reasonable time.

We may note at this point that collective burden-sharing does not necessarily imply that everybody formally shoulders the same share of burden. In practice, quite the opposite is true. Tax systems usually do not implement formal equality, but take more from those who have more. Accordingly, the costs related to benefits that are financed from this resource are borne collectively but cause disparate burdens on different individuals. The same applies, *mutatis mutandis*, to social insurance schemes. Here, too, redistribution often takes place not just on the benefit side, but also in relation to revenues because the level of contributions differs. So, formal inequality is typically involved also when the burden is shared collectively.

Consequently, the particular problem with AAPs is not the unequal distribution of the related burden *per se*. However, unlike AAPs, in the cases of collective burden-sharing referred to above, the uneven distribution is deliberately designed so as to take account of disparate capabilities of shouldering the burden across the polity. The idea, put simply, is that richer individuals carry higher burdens than poorer ones. There is, to be sure, reason to challenge this as an unrealistic idealisation. Probably, one could spot inconsistencies – and maybe even blatant injustices – in most countries' systems of collective burden-sharing. But at least in principle, such burden-sharing follows a distributional scheme, however imperfect it may be in practice.

AAPs, by contrast, do not appear to implement any systematic scheme of burden-sharing, at least not at first sight. Rather, they seem to place a significant burden on only a few individuals. At the same time, it is open to question whether these burdened individuals are peculiarly well equipped to shoulder this burden – hence the suspicion that such unequal treatment could be unjustified, and hence the related objection specifically to AAPs.

Thus, the key question is whether and how the disparate and seemingly unsystematic burdening of certain individuals that is characteristic of AAPs can be justified. Finding such a justification, or at least justificatory strategies, is essential to our overall endeavour of weakening the case against AAPs. In order to do so, we point first to similar patterns in other regulations (3.4.1.1), in a second step, we identify the respective justifications for their existence (3.4.1.2), and third, we discuss whether these can be transferred to our case of AAPs (3.4.1.3).

3.4.1.1 Burdened individuals in other contexts

We have shown that collective – and systematic – burden-sharing is common in EOPs. However, this observation does not always apply. In fact, we can find exceptions in different branches of the law.

3.4.1.1.1 Burdened individuals in employment law

Many of these exceptions belong to the field of employment law. Think, for instance, of rules concerning maternity or sick pay, or of minimum wage regulation. In these

cases, as in ours, the respective policy aims at promoting equality – here, it is a kind of distributive equality for sick employees, mothers (-to-be), and low-paid workers. And in all these cases, it is not the entire polity, but the employer who bears the (immediate) burden.

3.4.1.1.2 *Burdened individuals in family law*

The pattern we can discern in employment law is not confined to this field. We find the same structure in family law when dealing with duties of maintenance. Here again, individual persons are obliged to maintain other individuals, typically their relatives or current or former spouses. This is a goal which otherwise the entire polity would pursue, at least to the extent that we are dealing with minimum subsistence and – in some countries – deserving individuals. From this perspective, maintenance law appears as a substitute for policy instruments that are partially equivalent and would implement a collective burden-sharing scheme.

3.4.1.1.3 *Burdened individuals in the enforcement of judgments*

Another example is the law relating to the enforcement of judgments. Here, we find rules that prevent the creditor from recovering a debt to the extent that the debtor would then become destitute. The aim is to prevent individual impoverishment (which might in turn lead to a subjective entitlement to governmental support), but the burden is placed on the shoulders of the individual creditor.

3.4.1.1.4 *Burdening individuals – an exception?*

This list of exceptions could be extended. But it is already long enough so that we might ask whether it is appropriate that we speak of this pattern as an exception. In fact, the reason for calling it an exception is not so much quantitative as it is systematic. The underlying idea seems to be one of the fundamental individualist premises of liberal polities – namely, the separation between a public and a private sphere, with the latter being the realm of communicative and the former that of distributive justice.¹⁵

This does not mean, on the level of individual morality, that one person is not responsible for another. But it implies that one can, when confronted with the neediness of another person, legitimately point to the primary responsibility of the polity as a whole rather than accepting primary responsibility oneself. The view may

15 For this widespread notion based on Aristotelean categories, see e.g. Gustav Radbruch, *Rechtsphilosophie*, 2nd ed. (Heidelberg: C.F. Müller, 2003), 120; for more recent and differentiated discussions see Claus-Wilhelm Canaris, *Die Bedeutung der iustitia distributiva im deutschen Vertragsrecht* (München: C.H. Beck, 1997), 33 ff. Olha O Cherednychenko, *Fundamental rights, contract law and the protection of the weaker party* (München: Sellier, 2007), 44; and Christian Helmrich, *Mindestlohn zur Existenzsicherung*, dissertational manuscript on file with the authors, 190; the dissertation is forthcoming with Nomos in 2014.

well be contestable, especially if one lives in a system that is less than perfectly just.¹⁶ But it seems to correspond quite well with how people behave (and believe they may legitimately behave) in their daily lives. Imagine meeting a beggar on the street. Some people might be willing to give, but very few would accept a duty to do so, even if they thought that the beggar should receive help.

The reason is, and we are back to the objections to AAPs, that the burden on the helper would seem coincidental. Others might never encounter the beggar and thus go unburdened. Or they might be much richer, but not burdened according to their increased capability of helping. It is the issue of disparate burdening, or of the absence of systematic burden-sharing that these cases have in common. What is at stake, hence, is a violation of the equality principle.¹⁷ This is why exceptions to the rule of systematic burden-sharing require a justification – no matter how frequent they are.

The long list of examples, therefore, does not serve to remove the requirement of a justification. But it might indicate what kinds of justification are acceptable. This is the next section's focus.

3.4.1.2 Justifications for burdening individual actors

We have shown what the exceptions and AAPs have in common. However, the reasons that might justify these exceptions are different. In the following discussion, we will review these justifications and consider to what extent they may be transferred to AAPs and the two categories of burdened individuals.

3.4.1.2.1 Justifications based on the assumption of a moral responsibility

Maybe the most straightforward way of justifying an exception is to argue that there is a moral duty for the burdened individual to care for the beneficiary. And the most obvious case in point is maintenance law. Here, it is indeed likely to be widely accepted that, say, parents owe such a duty to their children, that this duty should be prior to any such obligation of the general public, and that it should be implemented by law. The sources of such a moral obligation may, on a closer look, differ. Among spouses, the voluntary act of marriage is probably the most plausible source of such an obligation nowadays. Similarly, the duty of parents towards their children could be viewed as arising from the deliberate act of procreation. But in practice, deliberateness may in some cases be doubtful, so one might consider replacing it with kinship as a source of such obligation. The opposite case (the duty

16 For an extensive treatment of the issue see Gerald A Cohen, *If You're an Egalitarian, How Come You Are So Rich?* (Cambridge: Harvard University Press, 2000), 148 ff.

17 For a thorough treatment of the issue and a plea for dealing with it primarily from the perspective of equality, see Christian Helmrich, *Mindestlohn zur Existenzsicherung*, dissertational manuscript on file with the authors, 190; the dissertation is forthcoming with *Nomos* in 2014; Part IV.

of children towards their parents) – provided that it is acknowledged at all – is also complex. Again, kinship is a possible explanation, but so is (some abstract notion of) reciprocity. In short, the exact line of reasoning might differ. But, in all these cases, the personal relationship is the reason for ascribing an increased obligation among individuals and which legitimises a policy that benefits one and burdens the other.

This pattern need not be confined to family relationships, but might be found also in other contexts. Employment law is a case in point. In this field, contracts typically come with an increased degree of permanence and maybe personal proximity between the parties. Hence, the justification for burdening an employer with sick or maternity pay, for example, may be sought in a moral duty of (arguably reciprocal) care and loyalty between employer and employee.

This justification is certainly more contestable here than in the context of family law. But it is still not implausible, especially not in legal systems whose employment law recognises such mutual duties. However, it does not necessarily work for all relevant examples from employment law. Minimum wage regulation, for instance, is harder to justify in this way because it applies already at the beginning of any employment contract, when there is no pre-existing relationship upon which the assumption of such an increased obligation could be based.¹⁸

3.4.1.2.2 Justifications based on practical considerations

Minimum wage regulation is not the only case in which the above line of reasoning does not work. The same is true for other contractual relationships without a comparable degree of permanence and proximity. From our list of examples, enforcement law comes to mind. It would hardly be plausible to claim that any creditor bears an increased moral responsibility for the debtor.

So, we need to look for another kind of justification. Here, and in other cases, it seems to be not moral, but practical considerations that account for burdening the private individual. Burdening the creditor is necessary to prevent a circle in which whatever the creditor takes from the debtor will be substituted by welfare grants that the destitute debtor might receive. And it might be justifiable because it only increases a risk that the creditor bears anyway – that is, the risk of the debtor's insolvency.

Similarly, minimum wage rules may be justified on pragmatic grounds. The same purpose, to be sure, could be achieved by relying solely on tax-financed benefits. But a plausible rationale that may justify burdening the employer instead is that wage regulation gives an incentive to work and prevents employers from keeping wages artificially low because of supplementary public assistance to the poor.

18 *Ibid.*

3.4.1.2.3 Assessing the relevance of these justifications

The considerations that we have sketched here may or may not be viewed as sufficient justifications. One might disagree with the moral ascriptions of responsibility as suggested before, and in the case of merely practical justifications, it might be all the more contestable whether they are important enough to outweigh the unequal burden to the individual. Accordingly, it is not just the kind and weight of the justification that is important. It is also relevant to identify on whom the burden is placed and how intense it is.

The above examples differ significantly with regard to these aspects. Most notably, we have to ask especially in the employment context whether the burden could possibly be passed on so it would ultimately be borne by consumers and hence spread much more widely. Similarly, one should recall that *de iure* we often treat entities as individuals which *de facto* they are not. The 'employer' will often be a corporation (which is why it seems more appropriate in such cases to speak of individual actors rather than individuals or individual persons). This, too, entails that the burden is spread. And even if either type of burden-spreading is not systematic in the way it is true for the collective burden-spreading schemes described before, both types do reduce the impact of the unequal burden on each individual person involved.

3.4.1.2.4 Interim summary

We end our sketchy review of potential justifications here. We have raised many issues and, admittedly, left most of them unresolved. However, it is not only impossible to solve any of them in the abstract setting of this article, without reference, that is, to any specific positive legal order, but it is also unnecessary for the purposes of our argument. We had shown before that it is not unusual for policies to place an unsystematic and disparate burden on individuals, but that such unequal treatment requires a justification, and we have now presented a variety of considerations that are relevant when considering the justifiability of such disparate burdens and of AAPs in particular.

3.4.1.3 Transferring justifications

AAPs differ from the policies we reviewed in the previous sections in that we can identify not just one, but two distinct individual actors who incur a disparate burden. These are the burdened addressee and the disadvantaged competitor. Any justification for AAPs will have to distinguish between their respective situations. We will do that when we now apply the above considerations to AAPs.

3.4.1.3.1 The unavailability of individual moral obligations as justification for AAPs

In the context of AAPs, it is hard to postulate a moral obligation on the burdened individual actors towards the beneficiaries. At least, personal proximity is unlikely

to play a role. This is evident in the case of the disadvantaged competitor who will typically not even (get to) know the beneficiary. The same is true for the burdened addressee if the AAP is targeted at the first contact between burdened addressee and beneficiary. This is likely to be the case for admission, appointment, or hiring scenarios. A counter example is the case of promotion decisions within one enterprise when there is indeed a pre-existing relationship. However, such situations will be less frequent, and they are also ambivalent because the employer will then be similarly proximate to the disadvantaged competitors, resulting in a collision of responsibilities.

A basis for assuming a moral obligation that is more plausible than personal proximity may be past injustice. AAPs are employed against the background of longstanding discrimination, and they might hence be interpreted as an attempt to rectify the disadvantages that their group of beneficiaries suffered in the past and that have persisted into the present. Today's benefits may hence be seen as a means to compensate for yesterday's injustice. The problem with this reasoning, however, is that the individual link across time cannot typically be established. AAPs cannot ensure that their beneficiaries have indeed suffered any disadvantage, nor that the burdened individual actors (disadvantaged competitor and burdened addressee) have enjoyed any advantage or caused any disadvantage.

This is not to say that such a personal link is indispensable to justify AAPs. We acknowledge that this issue is contested and reserve judgment on it. But we maintain that such a link would indeed be necessary if the justification for AAPs were to be based on a moral obligation on the burdened individual actors towards the beneficiaries to compensate them for injustice suffered in the past.

3.4.1.3.2 Justifying AAPs through practical considerations

While we cannot see how individual moral obligation might justify the disparate burden that AAPs entail, we think that practical considerations may be more successful in this regard. AAPs differ from the other examples we have listed in that their aim is usually not purely pecuniary. Of course, getting a job or a university degree may have positive pecuniary effects for the beneficiary. But that is not the only purpose of AAPs, and arguably not even the primary one.

Rather, they guarantee access to and representation in spheres of social life that would otherwise be foreclosed to many of their beneficiaries. This additional, non-pecuniary purpose renders them harder to replace. While in the case of pecuniary purposes there is always the alternative of employing a policy financed through a systematic burden-sharing scheme instead, this is not possible for AAPs.

If you want, for instance, more members of a disadvantaged group in tertiary education, it is not possible to just 'buy admission' for them instead of using an AAP.

There may be other options, to be sure, such as investing in their primary and secondary education so they may eventually not need an AAP. This could indeed lead to the same outcome, but only in the long-term and with less certainty – which illustrates our point. When dealing with AAPs, it is harder to find policies that are functionally equivalent than it is for policies whose purpose is pecuniary.

This is an important aspect, and it makes the justification of AAPs easier. It does not mean, though, that they are always justified. It may well be that their aim is not (considered) important enough to outweigh the disparate burden they entail. One needs to look, hence, to the other side and assess the weight of that burden. This weight will differ considerably for the two categories of burdened individual actors.

For the burdened addressee, there are various factors that may reduce this burden. First, the burdened addressee may not get the candidate of its preference. But it does get another candidate instead. As pointed out before, it will depend on the circumstances of the specific case – and on the design of the AAP – how relevant that burden is. Second, it may well be that this burden is spread on many shoulders – namely, when the burdened addressee is a corporation consisting of many individuals or when it is able to pass the burden on to its ‘customers’.

For the disadvantaged competitor, by contrast, none of these mitigating factors comes into play. She does not get anything instead, and must alone bear the burden, which may, as we have pointed out before, be significant.

3.4.1.4 Interim summary

Our discussion of the ‘argument from equality’ has shown that the unsystematic disparate burdening of individual actors that is characteristic of AAPs is by no means uncommon for policies that seek to promote equality. Nonetheless, such burdening requires a justification. This will, in turn, ultimately depend on the specific policy and the legal system in question. But it could be shown in our discussion above that the absence of alternative policy options that are functionally equivalent will typically serve as a relatively strong practical consideration in favour of AAPs, while the countervailing objections are different for either category of burdened individual actor. The disparate treatment that the burdened addressee suffers is mitigated by a number of factors and is thus easier to justify. However, this is not the case for disadvantaged competitors. It is here in particular that our defence of AAPs still needs to be expanded. We will come back to this aspect in particular once we have discussed the ‘argument from personal autonomy’ in the next section.

3.4.2 The argument from personal autonomy

As mentioned in the first section, there is yet a second critical objection directed against AAPs. Unlike the ‘argument from equality’, this second objection is not

concerned with the unequal distribution of the burden resulting from AAPs, but rather with the very nature of such burden. The perceived unfairness of AAPs stems in this case from the fact that the burden they create is per se problematic and should, at least in principle, be eschewed. The burden at hand refers to the unwarranted intervention in the personal autonomy of the addressees of AAPs who will have their decisions at least in part determined by these policies. The central claim of the 'argument from autonomy', then, is that AAPs unduly infringe upon the individual autonomy of their addressees.

Our discussion of this argument will begin by specifying the reasons that individual autonomy, according to common paradigms, is considered worthy of protection, and how this relates to our context (3.4.2.1). In a second step, we shall call into question the core premises of these paradigms. In order to do so, we will briefly introduce findings from research in social psychology (3.4.2.2), explain how they weaken the autonomy-based arguments in our context (3.4.2.3), and discuss their implications for the choice of EOPs in particular (3.4.2.4).

3.4.2.1 Why personal autonomy may be worth protecting

Infringement of personal autonomy is problematic for different reasons. One may, first, consider autonomous decision-taking an element of individual liberty. Any restriction will then be undesirable for the sake of the affected individual. Autonomy, in this view, is an end in itself.

But it may also, second, be considered a means to another, collective, end, and that is economic efficiency. Autonomy, in this perspective, is not only a source of personal well-being, but also a means of attaining an optimal allocation of scarce economic resources in society. Individuals in economic environments are taken to have the goal of maximising the utility they derive from the allocation of scarce resources, and therefore a rational and self-serving individual makes choices that she considers most 'efficient' – that is, suitable to attain the desired ends with the minimum amount of inputs. The aggregation of such efficient individual choices leads to collective utility maximisation. External interference in decision-making – inherent to AAPs – reduces such efficiency. Employers, for instance, are not allowed to choose the job candidates they consider best. Universities are not allowed to enrol the students whom they find most suitable. The efficiency of utility-maximising decision-making is impaired, and all of society is negatively affected through aggregate utility losses.

As these examples illustrate, the second, efficiency-related argument is but the flip-side of what we have discussed before within the context of the 'argument from equality'.¹⁹ This is because the alleged collective loss consists of the aggregated losses that the burdened addressees of an AAP incur by being forced to take (what they

¹⁹ See 3.4.1.

perceive as)²⁰ sub-optimal decisions. This loss is quantifiable in pecuniary terms. As EOPs often come with such a burden, AAPs face particular objections only to the extent that the individual is burdened disparately. This aspect is what we have discussed in the preceding section. It may therefore be disregarded here.

In our context, autonomy is therefore relevant only from the first perspective that views it as a component of individual liberty. This is not to say it is not important. Quite to the contrary, liberty is among the core values that modern polities commonly subscribe to, and individual autonomy no doubt is an important component of it, arguably even its most fundamental one.

It should be noted, however, that in our specific context the actors subject to such restrictions of their liberty are often not individuals, but corporate entities, such as universities and firms that act as employers. It is true, to be sure, that it is ultimately individual persons who act for such collective bodies. But the fact that they do not act for themselves, but for and as parts of such bodies may weaken the concern for their autonomy. The efficiency effects will be no different. But as we are focusing on the dimension of personal liberty, the difference would seem relevant. Corporations are not as such capable of personal liberty, and while the individuals who act for them are, it is quite common that the law requires these individuals to accept limitations of their liberty as long as they act for a corporation.

In sum, this section was meant to explicate why personal autonomy is worthy of protection – and to what extent this is relevant in our context. We have thus identified the object now of the ‘argument from personal autonomy’. In the next section, we will take a closer look at its underlying premises.

3.4.2.2 The notion of autonomy and its behavioural premises

The argument from autonomy is rooted not only in philosophical, but also behavioural assumptions. The latter are typically left implicit or underdeveloped in policy and legal debates. Both the value of individual autonomy and the presumed economic efficiency that results from it are based on the premise that individuals are capable of such autonomy – that is, that they can choose the goals to pursue as well as the reasons for doing so, and that they can decide and act accordingly.

Over the last decades, however, research from the field of implicit social cognition (ISC) has demonstrated that such behavioural assumptions rest on empirically shaky grounds. The idea of a rational and self-governing mind (the very behavioural underpinning of personal autonomy as a value and principle) has suffered severe blows from empirical research that reveals how little deliberate control and intent are actually involved in many of our most fundamental cognitive processes.

20 On the relevance of the difference between perceived and real loss in this context see 3.4.2.4.

This field of research is of crucial significance for understanding the causes and nature of inequality in contemporary societies. If its findings are reliable – and the depth and range of empirical evidence suggests they are²¹ – then individuals constantly stereotype, evaluate and discriminate without typically being aware of doing so. Beyond this – and even more problematic for the notion of a self-governing and ‘autonomous’ mind – such implicit constructions and processes often contradict the consciously endorsed values of actors and operate *against* their deliberate intentions, values and goals.

3.4.2.2.1 *The implicit dimensions of bias and behaviour*

Scientific interest in the unconscious dimensions of cognition and behaviour has existed since the development of psychoanalytic theory in the late nineteenth century. Nevertheless, it was not until the 1970s that a systematic and empirically oriented research programme emerged to study the unconscious mechanisms underlying mental processes and action.²² It was at this time that social psychologists, drawing from methods developed in cognitive psychology, began to discover that people’s actual judgements and behaviour systematically diverged from their consciously endorsed attitudes and beliefs. In seeking to understand these discrepancies, researchers found that in certain contexts, attitudes and behaviour could be better explained by mental processes that were neither controlled nor known by the individual.²³ These were referred to as automatic and implicit.

-
- 21 Anthony G Greenwald & Linda Hamilton Krieger, ‘Implicit Bias: Scientific Foundations’, *Cal. L. Rev* 94 (2006): 945, 946. For partial reviews of the emergence and consolidation of research in implicit social cognition, see: Anthony Greenwald, ‘New Look 3: Unconscious Cognition Reclaimed’, *Am. Psychol.* 47 (1992): 766, on early empirical evidence and measurement of implicit cognition; Nilanjana Dasgupta, ‘Implicit Ingroup Favoritism, Outgroup Favoritism, and Their Behavioral Manifestations’, *Soc. Jus. Research* 17 (2004): 143, on the specific development of research on implicit ingroup and outgroup preferences and their impacts on social interaction; and Irene V Blair, ‘Implicit Stereotypes and Prejudice’, in *Cognitive Social Psychology: the Princeton Symposium on the Legacy and Future of Social Cognition*, ed. Gordon B Moskowitz (Hillsdale: Erlbaum, 2001), 359. On the methodological validity and political relevance of implicit bias research, see John T Jost, Laurie A Rudman, Irene Blair *et al.*, ‘The Existence of Implicit Bias is Beyond Reasonable Doubt: A Refutation of Ideological and Methodological Objections and executive Summary of Ten Studies that no Manager Should Ignore’, *Research in Organizational Behavior* 29 (2009): 39-69.
- 22 Patricia Devine & Margo Monteith, ‘Automaticity and control in stereotyping’, in *Dual-process theories in social psychology*, eds. Shelly Chaiken & Yaacov Trope (New York: Guilford Press, 1999), 339.
- 23 For one of the first demonstrations of implicit intergroup bias, see Samuel L Gaertner & John P McLaughlin, ‘Racial Stereotypes: Associations and Ascriptions of Positive and Negative Characteristics’, *Soc. Psychol. Q.* 46 (1983): 23. The authors sought to assess the

Social psychologists use these terms to designate mental processes that operate outside the realm of conscious awareness and control, distinguishing them from their explicit and controlled counterparts. The terms 'implicit' and 'explicit' refer to the perceiver's level of awareness of a psychological process. A process is explicit if it can be consciously detected and reported (regardless of whether it was triggered spontaneously), and implicit if it cannot be directly inferred through introspective awareness.²⁴ The terms automatic and controlled are used to indicate whether a process is intentional – that is, desired by the perceiver – or unintentional.²⁵

3.4.2.2 *Measuring implicit bias*

But what exactly are implicit mental processes and how can they be measured? Obtaining access to mental constructs of which subjects are unaware posed important methodological challenges for empirical researchers. In fact, the development of a research programme on implicit cognition is directly linked to the development of new empirical tools that allowed scientists to access and measure processes that would, under normal conditions, be distorted by intentional dissimulation or concealed by introspective inaccessibility. Such 'indirect' (that is, unobtrusive) measurement techniques used devices such as time constraints and subliminally activated cues in order to access and measure mental processes that operated automatically (that is, without the subject being able to exert control) or implicitly (that is, without the subject's awareness).

strength of implicit stereotypes by measuring how quickly and easily certain traits and attributes 'pop into mind' when people see the name or picture of a particular social group. Their hypothesis was that implicit stereotypes should facilitate stereotype-congruent associations, and that the strength of such associations could be measured by the time it took individuals to identify stereotype-congruent words after being exposed to a cue of a social group. The researchers found that participants categorised African-American stereotype words more quickly when they were paired with the group label 'black' than the label 'white'. Respondents were faster to identify word pairs such as white-ambitious and black-lazy, and this was taken as evidence that the prime (social category) and target word (attribute) were included among common semantic networks that operate without the perceivers' awareness or control. Furthermore, the degree of bias on priming tasks was often unrelated to subjects' self-reported racial attitudes and beliefs. See also Patricia Devine, 'Stereotypes: their Automatic and Controlled Components', *J. Personality and Soc. Psychol.* 56 (1989): 5, on the dissociation between the implicit and explicit components of stereotyping and their impact on behaviour.

- 24 Anthony G Greenwald & Mahzarin R Banaji, 'Implicit Social Cognition: Attitudes, Self-Esteem, and Stereotypes', *Psychol. Rev.* 102, no. 1 (1995): 4–5. See also Timothy D Wilson, Samuel Lindsey & Tonya Y Schooler, 'A Model of Dual Attitudes', *Psychol. Rev.* 107 (2000): 101, for an account of the dissociation between implicit and explicit attitudes, and on the persistence of implicit attitudes even after explicit attitudes have been changed.
- 25 It is important to state that the dichotomy between implicit and explicit is a somewhat artificial one, and it would be more accurate to represent mental processes along a continuum that ranges from more implicit to more explicit – see e.g., Blair, 'Implicit Stereotypes and Prejudice', 361.

The development of the Implicit Association Test (IAT) in the late 1990s represented a crucial step in the implicit bias research programme. The IAT is 'a reaction-time measure that captures the strength with which social groups (and other attitude objects) are implicitly and automatically associated with good/bad evaluations and other characteristics'.²⁶ The test is used to assess strengths of associations between concepts by measuring response latencies in computer-administered categorisation tasks.²⁷ The implicit nature of the associations measured through the IAT is reinforced by the accelerated speed at which responses are required. The speed of the test is assumed to reduce or eliminate the interference of conscious awareness or control.

The use of such methods has allowed researchers to demonstrate that stereotypes and prejudice operate largely on implicit and automatic levels, meaning that individuals rely on stereotypes and prejudiced evaluations without realising or controlling these processes. The presence of implicit biases is widespread, and even consciously low-prejudiced individuals who subscribe to egalitarian values and condemn prejudice and discrimination have been shown to possess high levels of implicitly biased attitudes and judgements. And this is true for the vast majority of individuals. Bias against women, ethnic and cultural minorities, the elderly and other historically disadvantaged social groups is pervasive. A recent study compiling accumulated results of the IAT revealed that the majority of respondents had implicit preferences for young relative to old people, for European Americans relative to African Americans, for heterosexuals relative to homosexuals, and the list goes on.²⁸

26 See Jost *et al.*, 'The Existence of Implicit Bias is Beyond Reasonable Doubt: A Refutation of Ideological and Methodological Objections and executive Summary of Ten Studies that no Manager Should Ignore'

27 The IAT is designed in the following way: in an initial block of trials, exemplars of two contrasted categories (e.g. face images of young or old people) appear on the screen and subjects classify them by pressing one of two keys (e.g., 'i' for young people and 'e' for old people). Next, exemplars of another pair of contrasting concepts are also classified using the same keys (e.g., key 'i' for words with positive valence and 'e' for words with negative valence). Then, in a first combined task, participants categorise the four words using the two keys, each of which has two response options mapped into it (e.g., 'e' for young or positive and 'i' for old and negative). In a second combined task, the complimentary pairing is used (e.g., 'e' for young or negative and 'i' for old and positive). The difference in average response latency between these two sets is known as the *IAT effect*, meaning that larger IAT effects reflect stronger implicit associations between concept pairings. In the example above, faster responses for the [young-positive] and [old-negative] task than for the [young-negative] and [old-positive] task indicate a stronger association of young than of old with positive valence, and the greater the difference between average response latencies, the more the participant is biased towards the category 'young'.

28 Greenwald & Krieger, *supra* 'Implicit Bias: Scientific Foundations'.

It is worth pausing at this point to revisit our initial categorisation of different layers of EOPs in light of these findings. It may be hard to gauge exactly how relevant a factor individual discriminatory decisions are for social inequality – and how much need there is, hence, for third-layer EOPs that target this factor. But the pervasiveness of bias, as shown by these studies, no doubt underscores the relative weight of this particular cause of inequality, and of the importance of the policies that we are discussing in this article.

3.4.2.2.3 *From implicit bias to behaviour*

The implicit mental processes as described above would not be a problem for policymakers and legal practitioners if automatic and uncontrolled prejudices and stereotyping remained only in the minds of individuals and were never translated into biased responses and discriminatory conduct. But unfortunately, this does not seem to be the case. A growing body of evidence suggests that people do make significant social decisions and behave according to their implicit stereotypes and prejudices, without intending to or even being aware of doing so.

Furthermore, implicit biases have been shown to affect behaviour even in situations in which time constraints or attention manipulations do not pose major challenges to reflexive decision-making. A meta-analysis of 122 studies correlating implicit and explicit biases with behaviour has shown that in settings that involve inter-group interactions – that is, interactions between groups that differ in respect of ethnicity, age, skin colour or gender – the predictive validity of implicit bias measures significantly exceeds that of self-report measures. This data strongly suggests that the terms of such inter-groups interactions and discrimination are largely driven by implicit biases and not by explicit beliefs or attitudes. The majority of people judge and discriminate unknowingly, and this poses a definite challenge to how traditional ADL addresses this type of behaviour.

A further challenge stems from the fact that, in addition to being unaware of the biases that lead to discrimination, individuals have a very limited ability to correct these implicit mental processes. In studying how individuals reconcile their consciously endorsed egalitarian values with implicit stereotypes and prejudice, researchers have found that, when reminded of such values, actors are more likely to rationalise biased decisions than to reduce their level of bias.²⁹ This is done by reconstructing or adapting the criteria used in decision-making in a way that will justify the biased decision in a neutral matter. Such ex-post rationalisation process is termed casuistry,

29 See Michael I Norton, Joseph A Vandello & John M Darley, 'Casuistry and Social Category Bias', *J. Personality and Soc. Psychol.* 87 (2004): 817, for a general account of how, given the time and resources, individuals will tend to implicitly mask biased decision-making by deploying more acceptable criteria to justify their choices.

and it occurs independently of conscious awareness or control. This demonstrates that even when the individual has time and cognitive energy at her disposal, these resources may be used to rationalise bias instead of suppressing it.³⁰

Even more importantly for anti-discrimination policy, it has also been shown that certain social cues that increase pressure on the decision-maker to be objective will enhance casuistry instead of reducing bias. In a hiring-scenario experiment, Linder *et al.* have demonstrated that, when an equity norm prohibiting discrimination is enunciated during the decision-making process, participants retrospectively reported their decisions to be based on more bias-neutral criteria and less on social category information. Nevertheless, on average, the decision of participants remained just as biased in the presence of the equity norm as in its absence, showing that stating such a norm had the effect of increasing casuistry rather than reducing bias.³¹

3.4.2.3 Why personal autonomy may not be that worthy of protection after all

As we have seen in this section, the cognitive underpinnings of social interaction in general and of discriminatory behaviour in particular are difficult to reconcile with the notion of personal autonomy that underlies public policies in this field. We discriminate without even being aware of it – not all, to be sure, but most of us, and most of the time. What is more, even if made aware of it, we have only very limited capacity to control and correct our behaviour accordingly. And all of this applies regardless of, and often despite, our deliberate intentions and consciously endorsed values.

The question we are left with in light of this evidence, then, is: what autonomy are we talking about? Our mind does still appear to be capable of a situative will, of course. And it is plausible to label this mind ‘autonomous’ in the sense that it is not governed by anyone else (although there is growing evidence on its dependence upon external cues and thus on its susceptibility to manipulation).³² But does this

30 *Ibid.*

31 Nicole M Lindner and Brian A Nosek and Alexander Graser, ‘Age-Based Hiring Discrimination as a Function of Equity Norms and Self-Perceived Objectivity’ (31 May 2012), available at SSRN: <<http://ssrn.com/abstract=2071447>> or <<http://dx.doi.org/10.2139/ssrn.2071447>>.

32 For partial reviews of the literature on the malleability of implicit bias, see Irene V Blair, ‘The Malleability of Automatic Stereotypes and Prejudice’, *Personality and Soc. Psychol. Rev.* 6 (2002): 242; and Margo J Monteith & Corrine I Voils, ‘Exerting Control Over Prejudiced Responses’, in *Cognitive Social Psychology: The Princeton Symposium on the Legacy and Future of Social Cognition*, ed. Gordon B Moskowitz (Hillsdale: Erlbaum, 2001), 375. See Nilanjana Dasgupta & Anthony G Greenwald, ‘On the Malleability of Automatic Attitudes:

mind really deserve respect for its autonomy? Or rather, as our unitarian notion of the mind becomes increasingly doubtful, does the part of our mind that is particularly determinative of our social interactions deserve such respect, at least when dealing with matters of discrimination? And would it not, conversely, even be warranted to intervene in favour of the deliberate, conscious part of our minds?

It would seem that this conscious, deliberate will is what traditional notions of autonomy presume. If this presumption is eroded, so is the basis for affording autonomy the close-to-sacred position that it enjoys in our law and morality. And this is true for both perspectives from which this position can be justified – the efficiency and the individual liberty perspective.

This is not to say that we need to abandon the idea(1) of personal autonomy altogether. There are elements left of it, and the degree of its erosion depends on the specific context. But – and that is our core conclusion here – the reasons for protecting personal autonomy are particularly weak when we deal with the kinds of social interaction that are targeted by third-layer EOPs in general and by AAPs in particular.

3.4.2.4 What this may imply for public policy

What this means for the initial question of this section is evident. In light of the findings from ISC, the ‘argument from personal autonomy’ loses much of its weight. AAPs will hence be easier to justify in that regard.

This point can be reinforced if we consider not just AAPs and their absolute desirability, but widen the perspective so as to look at their relative attractiveness compared to other policies. More specifically, we pointed out before that the most widespread among those EOPs that target individual decisions – among third-layer EOPs, that is – are not AAPs, but ADLs. Further, we speculated that such prevalence of ADLs may be explicable by the fact that they are less intrusive with regard to the personal autonomy of their addressees than is the case with AAPs. Now, in light of the evidence presented above, such preference should be reconsidered. This is partly because, as we have just pointed out, any argument from personal autonomy becomes less powerful when assessed against this background. But it is also because ADLs appear much less promising. For how could a command not to discriminate achieve its goal if its addressee is not even aware of her bias, let alone able to control it? The implication is that the relative attractiveness of AAPs compared to ADL should increase.

Combating Automatic Prejudice with Images of Admired and Disliked Individuals’, *J. of Personality and Soc. Psychol.* 81, no. 5 (2001): 800, for evidence on the impact of exposure to counter-stereotypic group members on automatic prejudice.

Moreover, as we have said, our considerations on the argument from personal autonomy have an impact also on the argument from equality. We ended our discussion of that latter argument stating, that the disparate treatment that the burdened addressee suffers was typically easier to justify than that of the disadvantaged competitor because in the latter case, we could identify fewer mitigating factors.³³

Now, the findings presented in this section add such a mitigating factor for this case in that they cast new light on the term 'disadvantaged competitor'. Taking into account the data on the pervasiveness of implicit bias, these competitors are very likely to be advantaged on innumerable occasions in daily life. In many and maybe the majority of cases, the effect of AAPs may thus be described more properly as an attempt to level out advantage, rather than as bringing about disadvantage for these competitors.

Along the same lines, the pervasiveness of bias also entails that the perceived economic loss that the burdened addressee incurs from being prescribed a choice different from its first priority will in many cases not be real. This is because such preferences of the burdened addressee are likely to be affected by implicit bias and hence not to be governed exclusively by economic rationality. This in turn means that one cannot generally presume that the foregone choice would have been the best choice, and necessarily better for the burdened addressee than the one imposed upon the addressee by the AAP in question. In other words, given the quantitative dimensions of implicit bias, there are likely to be many cases in which an AAP does not constitute a 'burden' to its addressee, but rather helps her avoid a biased and economically unwise decision and replaces it with an economically wiser one.

3.5 Conclusion

Our conclusion comes as no surprise. We announced from the outset that we were making a case for AAPs – or rather 'the case against the case against AAPs', because our strategy was to identify the critical objections against AAPs and then to counter them. For this purpose, we contextualised AAPs within the broader spectrum of related policies, reviewed the potential objections briefly, distinguished between equality and autonomy-related ones, and discussed them from different disciplinary perspectives.

More specifically, regarding the 'argument from equality', we suggested that a core problem with AAPs is that the burden that they entail is mainly borne by certain individual actors, and we discussed how such unequal treatment can be justified. In

33 See 3.4.1.4.

this regard, we distinguished between two categories of individual actors, burdened addressees and disadvantaged competitors, and found that unequal treatment of the former category will typically be justifiable. With regard to the latter category, our preliminary conclusions were more ambivalent. Although the same justificatory considerations apply as for the other category, the burden seemed weightier here, and its justification hence more difficult. This initial assessment could be corrected, though, in light of our discussion of the 'argument from autonomy', an implication of which is that the burden is likely to be less relevant in many of these cases than is commonly perceived.

Regarding the 'argument from autonomy', we showed how the idea of an autonomous mind has suffered a hard blow from recent findings in the field of implicit social cognition. Individual decisions are constantly made based on criteria and premises that the individual is not aware of and cannot control, and that are systematically biased against disadvantaged social groups. These findings call into question the behavioural underpinnings that justify the protection of individual autonomy in various contexts, including ours in particular. It becomes clear against this background that the ideal of personal autonomy cannot justify the preference for other third-layer EOPs to the detriment of AAPs. Moreover, it also follows from these findings that with regard to the 'argument from equality', the unequal burden that a 'disadvantaged' competitor incurs due to an AAP may often prove not to be a burden at all, but rather may just level out the undue advantage that this competitor enjoys because of implicit biases on part of decision-takers.

In sum, the result of our endeavours has obviously not been to sweep away the critical objections to AAPs. It would have been naïve to try to do so after decades of learned debates on this matter. But we hope to contribute to sharpening the issues, and indeed also to weaken these objections considerably. So, in short, the game is not over. The justification of AAPs is still contestable. But we are playing the ball back to their opponents. Our article, hence, is an extensive 'Why not?'