

The Ontology of Discrimination

October 14, 2020

Abstract

Discrimination is a social phenomenon which seems to be widespread across different societies and cultures. Examples of discrimination concerning race, class, gender, and sexual orientation are not difficult to find in western contemporary societies. In this paper, I focus on the ontological ground of such phenomenon, with particular attention to its diffuse and institutionalised forms. I defend a broadly speaking reductionist approach, according to which the various manifestations of discrimination are grounded on the existence of the effects of “discriminatory stances” in social contexts. Discrimination may become part of the institutional sphere, either by way of bottom-up “crystallisation” of discriminatory practices, or by top-down “dilution” of institutional defaults into ordinary interactions.

Keywords: Discrimination · Social Ontology · Institutions

While the concept of discrimination has been investigated in political and ethical philosophy, it has received little to no attention within social ontology. Yet the question of which state of affairs in the social realm have to obtain in order for there to be discrimination in the world is both legitimate and interesting. It is legitimate, because it does not seem that we can trivially settle it by establishing the ontological status of social entities such as groups, organisations and institutions; and it is interesting, because the concept applies to different domains. In the ordinary manner of speaking, we use the vocabulary of discrimination to talk about individuals (“you shouldn’t do that. It’s discriminatory”), but also about organisations (“the council discriminates working class students”), institutions (“where same-sex marriage is not allowed, marriage is discriminatory”) and institutional entities such as rules, policies and laws (“the legislation is discriminatory towards ethnic minorities”). While attributing discriminatory behaviour to individuals is relatively unproblematic, the application of the concept to groups, institutions, organisations and rules is not straightforward.

If discriminating requires cognitive and conceptual capacities, as it seems plausible to maintain, then either groups, institutions and organisations possess such capacities, or discrimination talk in such cases is indirect, if not an outright metaphor. The same goes for discriminatory rules or directives. In this paper, I explore a third way, according to which groups and entities in the institutional sphere can *literally* discriminate, although the phenomenon of discrimination is at bottom a matter of individual behaviour. The basic ideas are two: (i) that systematic and diffuse discrimination does not require a collective will or collective actions of some sort, but only the presence of converging “discriminatory stances” in social contexts and (ii) that rules that are effective in social and institutional contexts can be discriminatory by *de facto* leading to discriminatory individual behaviour.

Although related, the debate in the literature about joint action on whether collective entities have mental states and perform collective actions has not direct bearing on the question whether the presence of facts about discrimination of social groups at the institutional level requires more than converging individual stances. It is an interesting, and open, question whether collective entities, such as organisations, firms, assemblies, and whole societies have cognitive and behavioral capacities and thus could discriminate in the individual sense.¹ But in the present context it is not a central question, because, if I am right, discrimination at the group and institutional level does not *require* more than converging individual stances. Facts about discrimination thus can populate the institutional sphere even if no collective agent discriminates.

In Sec.1, I will begin with few preliminary distinctions, which are rough and approximate, but suffice to single out institutions and organisations within social phenomena in a broader sense. In Sec.2, I will investigate the *individual* sense of discrimination that applies to merely social sphere. In Sec. 3, I will address the case of *diffuse* discrimination. In Sec. 4 I introduce the *rule-based* sense of discrimination that applies to the institutional sphere, but it is ultimately reducible to the individual one. In Sec. 5 the case of *institutionalised* discrimination.

¹Although many people think that collective entities cannot have mental states and perform collective actions (see Tuomela (2002)), not everybody agrees (see Gilbert (2006), and List and Pettit (2011)). The debate is clearly connected to that of the so called “ontological individualism” of the social sciences: are individuals and their interactions the basic “building blocks” of social phenomena, or should we consider sometimes at least collective and complex entities as explanatorily more fundamental? In this paper I won’t touch upon those broader issues. I will, however, argue that the presence of discrimination in society, even at the institutional level, does not force us to give up ontological individualism.

1 The merely social and the institutional aspects of social contexts

With “social context” I mean a concrete situation in which (human) individuals interact with each other. Typically, a social context encompasses two kinds of social elements: facts about the behaviour of the relevant components of the society — the *merely social elements* — and facts about rules, norms, policies and procedures that are *effective* for those components — the *institutional elements*. A rule is effective when it generates *deontic statuses* (viz. specific rights and duties) that are binding for the relevant individuals or groups. In the minimal sense in which I will use it, institutions are *social practices codified by systems of rules that are effective in a given social context*. I will use the label “institutional entities” for rules, norms, policies, enactments, deliberations, laws and the like. The sense of institution is “minimal” for two reasons. Firstly, it is compatible with different views in play in social sciences and philosophy.² Secondly, it is silent with respect to how the deontic statuses that are binding for all the relevant individuals are generated in the institutional contexts³, and it is under-specified with respect to how the binding effect come about⁴.

Institutional elements, institutional entities, institutions and institutional contexts do not exhaust the whole of the “institutional sphere” within social contexts. Certain institutions have a “body”, namely there is a group (or a system of groups) of people that play an important role in establishing, promulgating, changing, and interpreting the binding rules and policies that individuate it. I will call such groups (or systems of groups) *organisations*. Sometime the body of an institution encompasses all the people that are in one way or another involved in it (like in the case of various governmental and private organisations), but other times an institution has more members than its body: for instance, an university has a senate and various faculties and departments (a system of groups) as its body, but counts many more members (the students and various employees).⁵ In the rest of the paper, I will inquire how discrimination works

²For instance, it is compatible with Hindricks and Guala (2014)’s “ecumenical” approach to institutions as rule-in-equilibrium, as well as with the three main conceptions they individuate (the regulative-rule based one, the constitutive-rule based one, and the equilibrium-based one).

³I do have a specific (and to some, idiosyncratic) account about how deontic statuses are differently generated in the merely social and the institutional cases (see Torrenco (2016)), but since nothing crucial in my main points here hinges on it, I will leave it out of the picture.

⁴In particular, I do stay neutral on whether the rules are effective because they are generally complied with or because there is common knowledge of the sanctions one undergoes if brakes them. Thanks to Samuele Chilovi for discussion on this point.

⁵Note that here I am using organisation in a slightly broader sense than the ordinary sense, since the senate of an university or a department is not usually called an organisation. Thus, in ordinary language “organisation” has stricter uses than the one I use here. However, ordinary language also counts looser uses (which I won’t take into account here), as when we use it for a group of people

in the merely social sphere, and how it can reach into the institutional sphere.⁶

2 Discrimination in the individual sense

In this section, I restrict my discussion to the sense of discrimination that requires cognitive capacities and behavioural capacities, which I take to be the most unproblematic one. Given that individuals (i.e., persons) are prototypical bearers of such capacities, I will call it the *individual* sense of discrimination.

My analysis of discrimination in the individual sense distinguishes three related conceptions: the *neutral* conception, the *behavioural* conception, and the *moral* conception of discrimination. I characterize the neutral conception of discrimination as follows.

Discriminate₁ (Neutral conception): to individuate a plurality of persons in virtue of the way one takes them to differ from the others with respect to certain prominent properties.

Discriminate₁ is based upon a cognitive capacity, namely the capacity to individuate⁷ a plurality of persons.⁸ The persons that are individuated through an act of discrimination₁ do not need be the object of visual response or some other direct form of acquaintance. One can generically individuate persons by simply entertaining a general representation of them, as when someone thinks of the persons that usually live in a certain part of a city, or have a certain skin colour, or have certain sexual preferences.⁹

Given that discrimination₁ requires a contraposition between the persons having certain features and the *others*, that is the rest of the society, when discriminating₁

that is structured in some way, and perform some social function, even if it does not embody an institution—for instance the organisation for the annual new year bingo of the neighbourhood.

⁶Note that the contraposition between merely social elements and institutional elements does not map onto that between the *informal* and the *formal*. There certainly are many non-institutional practices that are very formal in nature. For instance, in many contemporary societies certain rules of behaviour concerning specific circumstances (e.g., exchanging business cards in Japan, dining at a club in Great Britain, etc.) are quite strictly regulated, and may even be written down in official documents.

⁷I am using ‘to individuate’ in a somewhat idiosyncratic way, that is why I specify its meaning in what follows.

⁸Here I am interested to a neutral conception of discrimination that is restricted to persons. There is probably an even broader neutral sense of ‘to discriminate’ in ordinary English, as in “I can discriminate green from red”.

⁹Note that neither the individuating properties nor the property attributed to the individuated people need to be *accurate*. For instance, one can discriminate₁ a group of people by representing them as the drug addicts usually loitering in a certain park, even if they are not drug addicts, and they have been (systematically) mistaken for the usual loiters of that park. Thanks to Esa Diaz Leon for having pushed me to be more clear on the accuracy point.

we are always singling out a *group* of people¹⁰. However, discrimination₁ does not necessarily involve *perceived membership to a group* as a individuating property, neither necessarily entails the recognition of certain features as typical of a certain group.¹¹ I will call that the group that is signalled out by discriminatory₁ behaviour the *target* of discrimination₁. Someone who is among the persons in the target of discrimination₁ is a *subject of discrimination₁*. This notion allows us to move from the neutral conception to the behavioural conception of discrimination.

Discriminate₂ (Behavioural conception): to treat someone in a certain way on the grounds of (i.e., in a manner that is sensitive to) being the subject of discrimination₁.

Even if the property on whose ground someone discriminates₁ need not be perceived membership to a group, it is important to notice that as a consequence of there being (possibly many) discriminatory₂ acts, it is always possible that a group becomes over time more and more prominent through different social contexts, and thus become a *salient social group*, that is a group that is “important to the structure of social interactions across a wide range of social contexts” (Lippert-Rasmussen (2006), 169). Thus, even if not relevant for the ontology of discrimination, the concept of *belonging to perceived social group* plays a crucial explanatory role in understanding what can be *morally wrong* with discrimination, given discriminatory situations can be instances of social injustice.¹²

The behavioural conception has moral import, since any description of a behaviour can be evaluated as the description of *morally justified* (i.e., good) behaviour or of a *morally unjustified* (i.e., bad) one.¹³ The fact that discrimination₂

¹⁰The group in question is a “social group” in the (minimal) sense that the property through which it is individuated is a *represented* or attributed property. However, it does not need to be a social group in the (more robust) sense that the property through which it is individuated is a socially relevant property.

¹¹Note also, a bit more pedantically, that to individuate a plurality of persons may not be the same as to individuate a *group* of persons. The first capacity requires some sort of plural reference, which may not involve any notion of group. The second capacity require mastering some notion of a group. I say that “may not be” rather than “it is not”, because if groups are to be identify with pluralities (Horden and Lopez de Sa (forthcoming)), then the two capacities may turn out to be the same. (Thanks to Dan Lopez De Sa for discussion on this point.)

¹²Cf. McGowan (2009) on the connection between discrimination and oppression. Thanks to Laura Caponetto and Bianca Cepollaro for useful discussion on this point.

¹³The behavioural conception allows also for *trivial* forms of discrimination. For instance, consider behaving in the following way: condemning acts that have been individuated as morally bad (More precisely, condemning people *when* they act in ways that are perceived as morally bad.) Such a behaviour would be a trivial form of discrimination₂ in the sense that by describing it we do not provide much information about the moral status of the behaviour. Surely, behaving in such a way would not be something that *per se* generates or is specifically connected to unjustified inequalities concerning this or that social group.

can have good and bad moral import should not be confused with the fact that it allows for positive and negative forms, as specified in the definition below.

Positive discrimination₂: having a certain positive behaviour (e.g. rewarding, advantaging) towards the discriminated₁ individuals.

Negative discrimination₂: having a certain negative behaviour (e.g. punishing, denying advantages) towards the discriminated₁ individuals.

Neither positive nor negative behavioural discrimination is *per se* morally good or morally bad. Negative behavioural discrimination can be morally justified. If someone gets fined because he was driving drunk has been the subject of an act of negative normative discrimination. However, to fight against the discrimination of the drunk drivers does not strike me as a particularly sound moral cause (and the reader can easily come out with another example, if they do not find this one convincing). Positive behavioural discrimination can be morally dubious or outright wrong. If redheads get all sorts of privileges at work because the boss has a soft spot for people with reddish hair, redheads in that work environment are subjects to positive behavioural discrimination. But such sort of discrimination should be morally (and possibly legally) condemned. What is morally relevant in discriminatory₂ behaviour is not whether the way the discriminated₁ group is treated is advantageous or not, but rather whether the behaviour itself is morally justified or not. This leads us to the moral conception of discrimination, which I define, in a first approximation, as follows.

Discriminate₃* (moral conception, first approximation): to discriminate₂ in a morally unjustified way.

For the purpose of this paper, it is not important to specify the concept of moral justification at issue here. Typically, discrimination₃* is a matter of individuating persons on the ground of *non-morally-loaded* categories (e.g., color of the skin, ethnical traits, sexual preferences), and as a consequence behave in a *morally significant* ways (e.g., disrespect, hostility). But more generally, the relevant notion of moral justification seems to involve a certain *appropriacy* between the discriminatory₂ behaviour and what discrimination₁ is based on. Consider again the example of the drunk driver. It is not discrimination₃* to fine the drunk driver, because the punishment is appropriate to the behaviour. In a society where drunken drivers are stripped of their vote rights, and incarcerated for life, it may be morally justified to fight against such inappropriate punishments, that is, against the discrimination₃* of the drunken drivers.¹⁴

¹⁴Note also that it is not the case that in all cases of unjustified discrimination one makes use of

The definition about is a first approximation for two reasons. Firstly, any morally loaded notion of discrimination should be contextually sensitive, while discriminate₃* is defined in an absolute way. Although cases of discrimination₂ (and ₁) requires social contexts to happen, they can be characterised and individuated *abstracting away* from social and institutional facts that does not pertain to the individual who discriminate₂. This is why we can define them without relativise them to social contexts. Contrariwise, it would be usually underdetermined whether an act of discrimination₂ is morally justified if we only look at what pertains to the act itself. Think of negative discrimination₂ against drivers without a driving license. In many contemporary societies it seems a morally justified form of discrimination, if the punishment is appropriate, for instance a fine. However, if we were in a context in which obtaining a driving license were very costly and did not track much the actual capacity of the drivers, but rather their social status, our moral evaluation of the discriminatory₂ behaviour would be probably different.

Secondly, given that discriminating₃ requires discriminating₂ and discriminating₁, and that by discriminating₂ and ₁ one always singles out a certain group (even though not necessarily a salient social group), to which a certain behaviour is sensitive, discriminatory₃ behaviour is a kind of morally relevant behaviour that involves in some sense a group K. Thus, a less approximate definition is the following.

Discriminate₃ relative to a social context C and a group K: to discriminate₂ in a way that in C is morally unjustified towards the K.

Now, it is a plausible empirical hypothesis that the “interesting” cases of discrimination₃ (the ones that are discussed in the literature and rightfully fought against by activists) are those with respect to which, because of systematic and diffuse discriminatory₃ behaviour a social injustice for people belonging to a salient social group is generated. However, it is important to stress though that when a salient social group K is discriminated₃, the facts that explain why the Ks are a salient social group are *not* part of the ground of the discrimination. It is rather the other way around: facts about the systematic and diffuse discrimination of the Ks may explain why the Ks are (or have become) a salient social group. In other words, it is the presence of people with discriminatory₃ behaviour—e.g., the homophobes, the racists, the misogynous, etc.—in given social contexts that creates discriminated₃ target groups.

morally loaded individuating category. For instance, during a famine, it may be morally *unjustified* not to give food to convicted criminals. Thanks to Francesco Guala for discussion on this point.

3 Diffuse discrimination

If the analysis of discrimination in the individual sense given above is on the right track, such form of discrimination does not require the existence of collective entities endowed with collective mental states and capable of collective behaviour. It does not require that on the side of the discriminated, since the targeted group may become a socially salient group after, and indeed in virtue of, the discriminatory₃ behaviour it is subject to. (From now on, unless otherwise specified, with “discrimination” and cognate terms I mean discrimination₃.) And it does not require that on the side of who discriminate, since the behaviour is ultimately individual behaviour. However, one may doubt that talking of diffuse discrimination and of contexts in which there are groups of people who systematically discriminate—e.g., the homophobes, the racists, the misogynous, etc.—requires *merely* an appeal to individual behaviour. In this section I will show that, even though in order to account for systematic and diffuse discrimination a more sophisticated story than the one so far told is needed, no collective capacities or collective actions shows up in such a story.

In most ordinary contexts in which a social group K is discriminated, the ground of the discrimination involves at some level the presence of a *common ground*. I take the idea of common ground from the analysis of linguistic communication and apply it to social interactions in general (Cf. Langton (2018)). Although in the analysis of linguistic communication the idea of a common ground is often exploited because of its dynamic aspects (things can be introduced and eliminated from the common ground via linguistic interactions), its static aspect is the most relevant for the study of discrimination. The static aspect of the common ground is constituted by *default options*. Those pertain not only to the factual dimension (roughly, what people in the context take to be true), but also (and more importantly here) the normative, evaluative, and emotional dimensions (See Langton (2012)). In particular, they concern things like:

1. Who has the authority to do certain things or allow others to act in certain ways
2. Who is to be blamed/appreciated for such and such behaviour
3. Who is the object of negative/positive feelings in a given context

The idea of a common ground that is useful for the study of discrimination is thus that of a set of cognitive, normative, evaluative, and emotional directives that have an *authoritative status*, in the sense that they are assumed (and known to be assumed) by the people involved in the social context to be the correct ones unless the circumstances prove otherwise.

Note that the elements in the common ground are *collectively* assumed by a certain social group (or, potentially, a whole society). It may be argued that collective sharing of this sort requires a kind of coordination that cannot be grounded merely in individual mental attitudes. If so, the reductionist slant of the present account of the ontology of discrimination seems in danger. Now, perhaps interaction at the level of a common ground requires an ontology that goes behind that of individuals and individual mental states. But that does not change the fact that the facts involving discrimination depends only on the effects of the individual stances on the discriminated people, independently from whether the individual stances may be part of more complex phenomena (such as the common ground) involving ontological categories of a collective nature. This is clear if we look at the *sources* that the default options of the common ground can have, which are both of a merely social and an institutional kind.

Among the merely social traits that can influence the common ground in a way that can easily lead to discrimination in the individual sense of a given social group K, we find phenomena like the following.

(I.a) **Implicit biases** that lead to an unfair behaviour w.r.t. Ks

(I.b) **Conscious but unspoken beliefs** that lead to an unfair behaviour w.r.t. Ks

(I.c) **Outspoken beliefs** that support and lead to an unfair behaviour w.r.t. Ks

Insofar as sources such as (i.a)-(i.c) and similar ones require (conscious or unconscious) acts of individuation that singles out a group K, the unfair behaviour required by them is individual discrimination as characterised above. When an individual is in a mental state such as (i.a)-(i.c) above, they can be said to have a *discriminatory stance* towards people belonging to K. A stance is a *disposition to reproduce a pattern of behaviour that tends to have certain effects*—in this case, the unfair treatment of K's members.¹⁵

Although stances are features of individual behaviour, one can say that the stances of distinct individuals in a social context *converge* when they tend to lead to the same effects. Two empirical hypotheses that strike me as plausible regarding converging stances are the following. Firstly, in a social context in which we find a comparatively high number of persons with converging discriminatory stances who interact with each other, those stances will tend to *reinforce* each other, in the sense that the negative effects of the discrimination of the Ks will be in general stronger than in contexts in which the number of people with discriminatory stances is lower and/or the interactions among them fewer.

¹⁵Cf. Blackburn (1993)'s notion of "stance" and its link with descriptive and normative beliefs. The notion is also related to that of *alief* in Gendler (2008). See also Camp (2013)'s notion of a "slurring perspective".

Secondly, converging discriminatory stances towards a group K in contexts in which they reinforce each other tend to become *shared*. A stance becomes shared in a social context when many individual patterns of behaviour tend to lead to the same effect (that is, it converges), *and* the relevant mental states underpinning the disposition of the individuals having the stance are *similar*. Again, notice that stances can be shared in a more substantive sense too, for instance in the sense that involves a collective agent (because it is required for an explanation of the coordinative behavior based on them). Yet, the behavior that constitute the facts about discrimination would still be individual behavior shared in the sense of convergent and similar. It may be that a plurality of racists have to constitute a social group in some very robust sense to constitute an institution that is permeated and inspired by racist principles, but what grounds the facts about discrimination of a certain racial group in such an institutional context would still be merely the effects of the behavior of the individuals (more on this in the next section).

Of course, even if we grant that it is sufficiently clear what are the relevant underpinning states here, the notion of similarity between them has to be strengthened up. Roughly, the idea is that the mental states in question have to be similar with respect to the *representational properties* that motivates them—that is, the mental states of people sharing a certain discriminatory stance have similar representational contents *that explain why* they have such stances. For instance, in contexts where we find many racists who discriminate who is thought to belong to a certain ethnic group K, the racists will typically have similar beliefs about the Ks (they are lazy, prone to crimes, violent, etc.) and similar emotional attitudes towards the Ks (they feel uneasy in their presence, they don't want to share public or private space with them, etc.).¹⁶ Shared discriminatory stances, such as those of the racists or the homophobes, come in varieties: they can encompass far reaching and detailed *stereotypes* to only sparse and vague clusters of features attributed to Ks, and they can encode different forms of emotional attitude, from willingness of differentiating oneself from “them” to outright contempt for the Ks or even their dehumanisation. What is important here is that besides the individuating properties (those on which discrimination₁ is based), other descriptive, emotional and evaluative elements attributed to the persons in the target group constitute a shared stance.¹⁷

In social contexts where discriminatory stances towards a social group K are

¹⁶One could try to strengthen up the notion by saying that they have to be similar with respect to the *informational states* they encode. However, it is not clear that denigratory stances carry information at all, given that they are often inaccurate with respect to the features they attribute to the Ks.

¹⁷Discriminatory stances have often a derogatory aspect. See for instance the role of stereotypes in the account of derogatory language in Jeshion (2013) and Williamson (2009). A problem which I am not discussing here is how discriminatory stances that plausibly lack propositional content, for instance implicit bias, can be shared. Maybe one can think of a further distinction between weakly

deeply entrenched and well-established and possibly publicly manifested by a non-insignificant part of the society, disadvantageous or hostile default options towards the K can easily enter the common ground and become a static part of it. For instance, homophobes and racists can be seen as groups of people with converging and possibly shared discriminatory stances, and if racist or homophobes are prominent across different social contexts, people perceived as belonging to the target group K become a social group to which disadvantageous or hostile default options are applied in many circumstances—that is, the social group K becomes salient in a negatively charged way. Regardless on whether the idea of a common ground does require collective mental capacities, diffuse discrimination can be explained through the notion of source of discriminatory elements in the common ground and of shared discriminatory stance, neither of which requires a beefed-up ontology with respect to the individual case.

4 Discrimination in the rule-based sense

As hinted at in a passage above, the sources of discriminatory elements in social interactions can have an institutional nature too. Whether marriage between same sex person is a permissible default option, for instance, typically depends on some institutional factor—such as whether a law that allows for it has been enforced. When in a social context a certain institution is said to discriminate or to be discriminatory, it is often hard to point out *which individuals* are discriminating. The attribution of discriminatory behaviour seems then to be directed towards the institution or organisation itself. In this and the next section, I argue that taking such talk at face value does not entail that the institution (or a related collective entity) possesses mental states, or performs collective actions. If so, institutions and organisations can literally (not metaphorically) discriminate, even if they do not discriminate in the individual sense.

The claim can be maintained by identifying a *rule-based* sense of discrimination that pertains to institutional entities such as rules, policies, and laws, and which can be defined in terms of discrimination in the individual sense, as follows.

A rule/policy/law is discriminatory (in a rule-based sense) for people of group K in a social context C in which it is effective iff in C it leads to

converging (that is converging in the sense I explain in the text, as mere convergence of the effects) and strongly converging, which is converging and sharing similar *underpinning cognitive mechanism*, even if those cannot be characterised in terms of representational content. Shared stances would then be a particular case of strongly converging stances. Thanks to Kourken Michaelian and Denis Perrin for discussion on this point.

discrimination₃ with respect to the Ks.

This rule-based sense of discrimination can be exploited to explicitly define a non-metaphoric way in which institutions and organisations can discriminate, as follows.

An institution discriminates *iff* the rules that codifies it are discriminatory

An organisation discriminates *iff* it is the body of a discriminatory institution

Although distinct, the *individual* sense in which individuals discriminate, possibly in a diffuse way through shared discriminatory stances, and the *rule-based* sense in which rules, institutions and organisations discriminate are strictly related. In both cases the presence of discrimination in the social context is grounded in behavioural facts concerning unfair treatment of a social group K. In other words, there is not a difference in ontological status between individual discrimination and rule-based discrimination. The difference between the two cases hinges rather on whether the unfair behaviour towards a group K originates from individual stances and attitudes or rules that bind and regulate individual behaviour.

In order to highlight the fact that discrimination in the rule-based sense does not require that a collective entity discriminate in an individual sense, and thus that a group possesses cognitive and behavioural capacities, two considerations are in place. Firstly, the intended meaning of “lead to” in the definition above does not back up the thesis that a collective intention is behind the issuing of the rule/policy/law. To a first approximation, a rule, in a given context, leads to discrimination towards the Ks if it gives a *substantive contribution* to the fact that in that context the Ks are discriminated. A substantive contribution is something in between possessing some role in the causal chain that brings about the discriminatory situation, and explicitly representing and prescribing a discriminatory situation. In other terms, a discriminatory rule does not need to prescribe to act in a (morally unjustified) disadvantageous way towards the Ks, but it is something *more* than a rule that allows for or is compatible with discriminatory behaviour. Rules that prescribes discrimination towards the Ks by explicitly identifying the social group K in their formulation are clear case of discriminatory rule, in any context in which they are effective. However, even if a rule does not contain an explicit prescription for a behaviour that is discriminatory towards the Ks, given a context, it can contribute to unfair treatment of a social group more than simply allowing for it, or having some causal role in it, and thus be discriminatory in this rule-based sense.

For instance, at the present, female athletes in Italy cannot be contracted by sport clubs as professional athletes, but only as amateurs. The negative (and unfair) effects of this situations are not merely symbolic: amateurs do not receive pension contributions or social security. This situation is not the outcome of a law that *explicitly* prohibits sport clubs from contracting women as professional athletes, but rather the outcome of a more complex legislative situation that leaves to each national sport association the right to select which *sportive* categories are entitled to contract athletes as professionals. Given that, as a matter of fact, for the sport associations it is advantageous to give such a right only to the most lucrative categories, no female category has ever been given such a right. The law in question counts as discriminatory towards women, in my account, even if it does not mention explicitly any gendered social group, and indeed in a very different social context it could, in principle, be either fair or discriminatory to some other social group.¹⁸

In connection with this first point, notice that there is in the literature a debate on the distinction between *direct* forms of discrimination and *indirect* forms of discrimination (see Eidelson (2015)), which should not be conflated with the distinction between individual and rule-based sense of discrimination discussed here. Roughly, direct forms of discrimination involve *intentional* behaviour based on the identification of someone as belonging to a targeted social group K, while indirect form of discrimination are situation in which *de facto* people belonging to K are treated in an unfair way by a law or a behaviour (as a consequence of more general social factors). If having intentions in some derivative sense (for instance, as expressed in a voting system) does not entail having cognitive abilities and mental states, groups and organizations could directly discriminate, even if they were not able to discriminate in the individual sense. More to the point, there presumably can be forms of direct discrimination in the rule-based sense, if an organization express an intention — for instance, when an assembly with legislative powers of some sort explicitly deliberates to treat in an unfair way a social group K. But there cannot be direct forms of discrimination pertaining to rules, policies, and laws as such. And this situation is independent from whether the rule, policy or law in question makes explicit reference to the social group K, or it is just a consequence of its application that the group K is treated unfairly. There cannot be form of direct rule-based discrimination because rules have no intention (even when they are

¹⁸As a matter of fact, in the actual context the law is discriminatory to many categories of male sport too, since at present only four sport categories can hire athletes at the professional level. It is still the case that *none* of the female category have a right to a contract for professional athletes— unless they are in the military. Further details here <https://www.iltascabile.com/societa/atlete-non-professioniste/>. A similar example: requirement of an ID for voting in US is often deemed to be discriminatory towards black people. Thanks to Adam Sennet for discussion on this point.

the outcome of intentional behaviour, or converging intention in a group).

Secondly, in the characterisation given above, a rule is discriminatory towards the Ks only with respect to a given social context and when it has certain effects on the Ks in such a context. Given that social contexts are usually characterised by whole systems of norms and policies, along with behavioural facts, discriminatory rules are always individuated in a somewhat *holistic* manner. Holism comes with a risk of overgeneralisation. In general, a rule that is useful for non-discriminatory purpose in a given context may still be part of a systems of rules that allow for discriminatory behaviour. For instance, the rule that prescribes to ride on the right allows for discrimination of black people, in any context in which black people are treated unfairly with respect to “stop and search” policies on the road. It would be strange, though, to charge the convention of driving on the right of being discriminatory. Besides, the holistic nature of the concept of rule-based discrimination is in tension with the idea that it is *specific rules* that leads, or give substantive contribution to discrimination. However, neither the risk of overgeneralisation nor such a tension should be overestimated. Even if the criterion of individuation is holistic, it seems to me that as a matter of fact we are more or less always in a position to identify the rules that more than others are responsible for the discrimination at issue. If this is correct, even though it may not be possible to have an a priori strategy to “filter out” the unwanted cases and to pick out exactly those rules that are *per se* discriminatory, we may be able to formulate *heuristics* to pinpoint the rules that contribute substantively to discrimination in a given context.¹⁹

5 Institutionalised discrimination

When one consider the institutional aspects of social contexts which are “toxic” for discriminated social groups, discriminatory stances can both influence in a bottom-up way the profile of the institutions we find in the social context—by way of what we may call a “crystallisation” of the discrimination, and be influenced in a top-down way by the institutional defaults we find in the social

¹⁹That does not mean that the same heuristic has to deliver satisfactory results in all cases, or that it will never be the case that more than one incompatible heuristics are eligible. Maybe in certain cases the heuristics will be trivial (because they will deliver the result that all rules are discriminatory in that context or none is), or empirically inadequate (because it is not trivial but delivers clearly wrong responses). Yet the framework I am sketching here does not require us to be *dogmatic* about heuristics. Heuristics can be refined, updated, and more generally different contexts may require different heuristics. Besides, the strategy to fight the discrimination that flows from a rule that is effective does not need to be in all cases to abolish or fight the rule in question. Maybe changing some of other rule that constitutes a condition for it to be discriminatory is the best strategy. (However, I grant that if this is systematically the case, the whole explanatory project I am suggesting here would funder.)

context—by what may be called the “dilution” of institutional discrimination into the merely social sphere. In this last session I will give an account of those phenomena and argue that in neither case the presence of institutionalised discrimination require us to “beef up” our ontological commitment with respect to the merely social case.

If we consider the common ground of social contexts in which discrimination in the rule-based sense is effective, we find sources like the following.

(II.a) Discriminatory rules that are entailed by **explicit institutional principles**

(II.b) Discriminatory rules that are entailed by **implicit institutional principles**

(II.c) Discriminatory **decisions** that are taken by organisations

Discriminatory rules can be among the explicit inspiring principles of an institution (case II.a). Those has been the case in the past in Europe, for instance, when fascist or nationalist parties have been in power for long enough to institutionalise their guiding ideas in a fully explicit manner. Case II.b is more insidious. Even when an institution has not explicit discriminatory principles, given that institutions come with “their own flavour”, there may be general, possibly unwritten tendencies that constitute the character or style of an institution (or for a certain period of time in that institution) and that systematically result in discrimination for a social group K.

In neither case, a discriminatory institution is involved in discriminatory behaviour in the individual sense. We do not need to assume that there is a sense in which institutions and organisations display an analogue of the cognitive capacities of individual persons in order for a social contexts to contain institutionalised discrimination. The ontological ground of this kind of discrimination is the discriminatory behaviour that follows by the fact that the discriminatory rules and principles are binding and effective in that social context. This fact does not rule out that the discriminatory institution may have emerged from the *crystallisation* of shared discriminatory stances into institutional principles. Indeed, there are several historical examples of explicit and implicit institutional principles that have raised from shared discriminatory stances. For instance, it is well documented how anti-semitic sentiments have been received and listened to by fascist and nazi governments before World War II (see Stanley (2015)). More generally, it is likely that people that are at present discriminated in a rule-based sense, for instance because of the laws of a certain state, are members of a social group that *has been* (and possibly stil are) the subject of discrimination in the individual sense by the individuals that supported the formation of that institution.

Besides, institutional elements and merely social one can either reinforce or interfere with one another. They may interfere because people may be (morally)

better or worse than the institutions “in” which they live. Thus, it may well be possible that an institutional principle is discriminatory in a given context only because of “institutional inertia” — that is, the only reason the principle or rule is still around is that it is difficult to change or modify, but almost no one in that context supports the content of the principle. In such a case its discriminatory effects can be widely mitigated by individual behaviours that aims at making the situation for the discriminated group as fair as possible. For instance, in a context where the institutional rules require that immigrants can receive medical assistance only if certain strict conditions are met unless the circumstances are extraordinary, doctors could *de facto* find ways to provide legally medical assistance by formally claiming extraordinary circumstances each time.

Conversely, institutions can be better than the individuals who interact in the corresponding institutional contexts, and individual biases and beliefs that would lead to actual discrimination (in the individual sense) may be counteracted by the effect of a fair institutional setting. For instance, even if the responsible of the hiring of an enterprise has beliefs (or biases) that lead them to prefer male candidates over female ones, the institutional policy of the enterprise may be such that it prevents the behaviour of the officer to be discriminatory by requiring that the CV’s of the applicants do not contain gender information. More problematic is the case in which the individual discrimination comes from someone having a crucial role in a organisation, as when officers refuse to register or celebrate same-sex marriages or unions in context where they are institutionalised.

Along with the crystallisation of discriminatory stances into institutional aspects, there is the converse phenomenon of the *dilution* of institutional elements into the merely social. Consider the case II.c above. Organisations such as firms, universities, and national agencies can enforce rules (or policies, laws, norms that entail them) that are, in a given context, discriminatory towards a given group K and they can do so by explicitly deliberating through the collective actions of their body. This is the case of a discriminatory rule being *decided* by a organisation. It is tempting here to take the talk of decision literally or quasi-literally. However, if we want to deny that, or at least remain agnostic on whether the concept of individual discrimination is applicable to organisations, the ontological ground of this kind of discrimination cannot be the “behaviour” of the organisation or the underlying institution. Indeed, even if collective entities can discriminate in the individual sense, there seems to be a kind of discrimination originated by the decision for which the alleged presence of a “collective mind” of the organisation as such is irrelevant. After all, it does not matter how the organisation has arrived to the decision to enforce a certain discriminatory policy or rule: maybe it was the outcome of a sincere unanimous vote, maybe the organisation has been swayed by toxicity of the so-

cial surrounding, or maybe it is just the success of the cunning subterfuges of an individual or a small group that lead — possibly through exploiting a proverbial useful idiot — to its official endorsement by the institution. Those factors are irrelevant, if the fact that the rule is effective leads to discrimination of a social group K in that social context, then the organisation is discriminating in the rule-based sense.

Both institutional principles and decisions can influence the common ground of the social contexts in which certain discriminatory institutions and organisations are present. This is so because, social interactions, even when informal, often involves *institutional defaults* — that is default options that have been established by institutional decisions or are the outcome of institutional principles. More generally, social interactions may rely on default options that are in the common ground in virtue of the presence, in the social context, of an institution or an organisation. Again, the presence of institutional default options do not require that the institutional sphere be endowed with collective faculties. Institutionalised discrimination is entirely grounded in discriminatory individual behaviour, even when its source is not the crystallisation of shared hostile attitudes into institutionalised elements, but rather the dilution of discriminatory principles and decision in the common ground of many social interactions, in the form of institutional defaults.

Acknowledgements

Besides the people referred to in the footnotes, and a helpful referee of this journal, I wish to thank all participants of the sessions of the *Seminar in Social Metaphysics* (University of Barcelona) and the *Laboratory's International Seminar* (University of Grenoble) where I presented drafts of a former version of this paper. For financial support ...

Giuliano Torrengo

University of Milan and Autonomous University of Barcelona

References

- Blackburn, Simon 1993, *Essays in Quasi-Realism*. New-York: Oxford University Press.
- Camp, Liz 2013. "Slurring Perspectives." *Analytic Philosophy* 54 (3), 330-349.
- Eidelson, Benjamin 2015. *Discrimination and Disrespect*. Oxford: Oxford University Press.

- Gendler, Tamar (2008) "Alief and belief", *Journal of Philosophy* 105 (10), 634-663.
- Gilbert, Margaret 2006. *A Theory of Political Obligation: Membership, Commitment, and the Bonds of Society*. Oxford: Oxford University Press.
- Hindricks, Frank and Guala, Francesco 2014. "Institutions, rules, and equilibria: A unified theory." *Journal of Institutional Economics* 11, 1-22.
- Horden, John and Lopez de Sa, Dan (forthcoming) "Groups as Pluralities", *Synthese*
- Jeshion, Robin 2013. "Slurs and Stereotype", *Analytic Philosophy* 54 (3), 314-329.
- List, Christian and Pettit, Philip 2011. *Group Agency. The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Langton, Rae 2012. "Beyond Belief: Pragmatics in Hate Speech and Pornography." In *Speech and Harm*. Edited by Ishani Maitra and Mary Kate McGowan, Oxford: Oxford University Press, 72-93.
- Langton, Rae 2018. "Blocking as Counter-Speech." In *Bad Words*. Edited by Sosa, David, New York: Oxford University Press, 144-64.
- Lippert-Rasmussen, Kasper 2006. "The Badness of Discrimination." *Ethical Theory and Moral Practice*, 9, 167-85.
- McGowan, Mary Kate 2009. "Oppressive Speech." *Australasian Journal of Philosophy* 87 (3), 389-407.
- Stanley, Jason 2015. *How Propaganda Works*. Princeton: Princeton University Press.
- Torrenco, Giuliano 2016. "Institutional Externalism." *Philosophy of the Social Sciences*, 47 (1), 67-85.
- Tuomela, Raimo 2002. *The Philosophy of Social Practice: A Collective Acceptance View*. Cambridge: Cambridge University Press.
- Williamson, Timothy 2009. "Reference, Inference, and the Semantics of Pejoratives." In *The Philosophy of David Kaplan*. Edited by Joseph Almog and Paolo Leonardi, Oxford: Oxford University Press, 137-59.