

Generics, race, and social perspectives

Patrick O'Donnell

Department of Philosophy, Johns Hopkins University, Baltimore, MD, USA

ABSTRACT

The project of this paper is to deliver a semantics for a broad subset of bare plural generics about racial kinds, a class which I will dub 'Type C generics.' Examples include 'Blacks are criminal' and 'Muslims are terrorists.' Type C generics have two interesting features. First, they link racial kinds with *socially perspectival predicates* (SPPs). SPPs lead interpreters to treat the relationship between kinds and predicates in generic constructions as nomic or non-accidental. Moreover, in computing their content, interpreters must make implicit reference to socially privileged perspectives which are treated as authoritative about whether a given object fits into the extension of the predicate. Such deference grants these authorities influence over both the conventional meaning of these terms and over the nature of the *objects* in the social ontology that these terms purport to describe, much the way a baseball umpire is authoritative over the meaning and metaphysics of 'strike'/'*strike*'. Second, terms like 'criminal' and 'terrorist' receive default *racialized* interpretations in which these terms conventionally token racial or ethnic identities. I show that neither of these features can be explained by Sarah-Jane Leslie's influential 'weak semantics' for generics, and show how my own 'socially perspectival semantics' fares better on both counts. Finally, I give an analysis of 'Blacks are criminal' which explores the semantic mechanisms that underlie default racialized interpretations.

ARTICLE HISTORY Received 10 November 2015; Accepted 1 March 2016

KEYWORDS Race; generics; social kinds; semantics; perspectivalism

Generics are generalizations about kinds such as:

- (1) Tigers are striped.
- (2) Birds fly.
- (3) Mosquitoes carry the West Nile Virus.

CONTACT Patrick O'Donnell  podonn14@jhu.edu

© 2017 Informa UK Limited, trading as Taylor & Francis Group

Despite their ostensible innocence, the problem of determining the truth conditions of generics has bedeviled semanticists for over 40 years. Part of the problem is that generics are not susceptible to any uniform quantificational analysis. For instance, despite the fact that (1)–(3) are each intuitively true and have the same surface structure, each sentence is true under different quantificational conditions. Sentence (1) seems to express something like *all tigers are striped* and thus requires a universal quantifier. On the other hand, (3) seems to be existentially quantified – it is true if there are at least some mosquitoes which carry West Nile. Finally, sentence (2) seems in need of a quantifier like ‘most.’ After all, most birds do fly, but penguins, ostriches, and over 40 other species of bird do not.

In fact, generic statements do not behave like kind-wide quantifications at all. For one thing, even paradigmatic ‘majority generics’ like (1) are not equivalent to any universally quantified sentence:

- (1) Tigers are striped.
- (1a) All tigers are striped.

Notably, (1), but not (1a), is compatible with the existence of non-striped tigers. Other, more tricky generics abound:

- (4) Ducks lay eggs.
- (5) Ducks are female.
- (6) Books are paperbacks.

Intuitively, (4) is true, and (5) is false. Why? After all, there are more female ducks than there are egg-laying ducks, since roughly half of ducks are female and a subset of that are unable or too immature to reproduce. (6) is intuitively false, despite the fact that 80% of all books are paperbacks. Thus, the sheer *prevalence* of a property F among a kind K does not seem to be evidence for the truth of a generic of the form *K's are F*.

One reason for this puzzle is that unlike quantifiers, generics express relationships between properties and *kinds*, not between properties and individuals. *How many* individual K's are F is frequently irrelevant to the truth-conditions of ‘K's are F.’ Generics also express special kinds of relationships between kinds and properties (Prasada and Dillingham 2006). They do not simply imply a statistical, accidental, or chance correlation between a kind and a property; rather, they imply – and are commonly interpreted as implying – a nomic, causal connection between the two. ‘Ravens

are black' does not mean that it just so happens that ravens tend to correlate with blackness at a rate of virtually 100%, but that something in the biological endowment of ravens makes it the case that they are black (and that the non-black ones represent mutations.) The same applies to weaker generalizations that do not invoke biological essences. Consider 'college students drink beer.' College students as a kind do not only drink beer at a statistically high rate, but there is something in the way that college life is set up (social camaraderie/pressure, access to alcohol, etc.) that suggests the presence of a causal connection between collegiate status and beer consumption. Generics thus express something over and above what is expressed by the kind-wide quantifications that are their cousins. '100% of/All ravens are black,' and '80% of/most college students drink beer' certainly imply strong relationships between ravens and blackness and college students and beer consumption, but they stop short of implying that this correlation is nomic or non-accidental.

Given their truth-conditional slipperiness, I am not optimistic that a semantics which accounts for *all* generics is possible.¹ Consequently, in this paper, I will not be exploring the semantics of generics *tout court*, but in pointing out some interesting features of a special class of generics that I will dub 'Type C Generics.' With a few notable exceptions (McConnell-Ginet 2003; Haslanger 2011, 2012, Leslie *forthcoming*), the generics literature has largely avoided this interesting and problematic class. Examples include:

¹This is not to say that ingenious and heroic global analyses have not been proposed. The three dominant approaches have been *normality* and/or *possible worlds* approaches (Asher and Morreau 1995; Nickel 2009), *stereotype* theories (Geurts 1985; Declerck 1986), and *probabilistic* approaches (Cohen 1999). Roughly speaking, normality approaches attempt to cash out the truth-conditions of generics in terms of some notion of normalcy. For instance, a generic statement 'As are F' might be interpreted as saying that *normal* As have property F, or that As do or would have property F in some set of *normal* possible worlds. Stereotype theories hold that 'As are F' is true iff it is the case that the *stereotypical* A has property F. Probabilistic approaches associate the generic truth-conditions of 'As are F' with comparative probabilities' which specify (1) the probability of an arbitrary x in A having property F or (2) the probability of an arbitrary x in A having property F as opposed to the probability of an arbitrary y in a superset of A having property F. Either strategy is employed depending on what the intuitive truth conditions of the generic statement are.

Each of these theories admits of intuitive counterexamples, and each makes at least some false empirical predictions. There is a lively literature which tries to address these a priori and empirical data points. See especially Carlson and Pelletier 1995.

- (7) Blacks² are criminal.
- (8) Muslims are terrorists.
- (9) Blacks are thuggish.

The project of this paper is to deliver a systematic treatment for these generics. The first goal will be to demonstrate that sentences like (7)–(9) comprise an interesting linguistic type. In so doing, I will discuss Leslie's (2007, 2008) influential 'weak semantics' for generics, and show that it does not account for the interesting features that this type has. Briefly, those interesting features are as follows. First, Type C generics link social kinds with special lexical items which I will call *socially perspectival predicates* (SPPs). Roughly, SPPs are special because they require interpreters to make reference to socially privileged perspectives which are authoritative about whether a given object fits into the extension of a given SPP. This capability is deserving of a special semantic treatment which Leslie's semantics does not offer. In particular, I argue that SPPs contain in their semantics a contextual parameter that is filled by a socially privileged *relevant perspective* (RP.) SPPs (and a fortiori, Type C generics) thus require what I will call a *perspectival semantics*.

Second, sentences like (7)–(9) trigger specific interpretations of the predicates attributed to them. These interpretations of SPPs like 'criminal,' 'thuggish,' and 'terrorist' are *racialized* – they link criminality, thuggishness, and terrorist activity to racial and ethnic identities, and imply that this link is nomic or non-accidental. In the final section of the paper, I give a perspectival semantics for 'criminal' which shows how these terms can be indexed to particular racial identities in interpretation.

1. Leslie's weak semantics: type A and type B generics

In her groundbreaking work on generics, Sarah-Jane Leslie has argued that the capacity to understand and interpret generics is a byproduct of a primitive generalization mechanism which is involved in early conceptual categorization and inductive reasoning. If Leslie is indeed correct that generics are indirect linguistic

²Throughout the paper I capitalize terms for races, as in 'White' and 'Black.' This is to distinguish race terms from color terms ('black' and 'white'), and also to draw attention to the fact that I am thinking of racial membership in terms of belonging to a certain social role, not only in terms of having a certain skin color.

evidence for the existence and function of such a mechanism, we ought to expect children's competence with generics to manifest itself fairly early in cognitive development.

A variety of empirical data lend *prima facie* support to Leslie's hypothesis. As early as 3 years old, children display adult-like competence with recognizing, interpreting, and reasoning with generic language (Gelman 2003). For instance, they seem to be sensitive to the ways in which the truth-conditions of generics depend on the precise nature of the property being attributed to the kind. Preschoolers will accept claims such as 'lions have manes' while rejecting claims such as 'lions are boys,' despite implicitly understanding that there are at least as many "boy" lions as there are maned lions.' (Brandone et al. 2012; Leslie 2014.)

On the other hand, preschoolers are far less competent when it comes to sentences involving quantifiers. Full competence with quantifying words such as 'all,' 'most,' and 'some' arises comparatively late in development (Gelman 2003; Halberda and Feigenson 2008) If we assume that computing more complex truth conditions involves more complex computations, this is a fairly surprising result. After all, the semantic profile of quantifiers seems to be far simpler and more uniform than the semantic profile of generics. Intuitively, 'All ravens are black' is satisfied just in case every raven in a given domain is black, 'most' is satisfied just in case more than half of all ravens are black, 'some' is satisfied just in case at least one raven is black, etc., whereas most semantic theories hold that the truth conditions of generics are rather complex. Surely we should expect comprehension of generics to piggyback on comprehension of quantifiers! Not only is this not the case, but it turns out that both children and adults *recall* quantified sentences as generics, suggesting that generics are both more cognitively fundamental and more easily interpreted than quantified sentences (Leslie and Gelman 2012.)

Despite the unruly semantic behavior of generics, adults and children seem to be in possession of a procedure for computing their truth-conditions, even when it comes to aforementioned 'tricky' generics like 'ducks lay eggs' or 'mosquitoes carry West Nile.' Leslie divides these tricky generics into two types, Type A and Type B.

Roughly, generics of Type A are intuitively true in virtue of (a) predicating properties which belong to a certain *characteristic dimension* for the kind and (b) predicating properties whose counter-instances are *negative* rather than positive. Plausible candidates for 'characteristic dimensions' include body morphology, reproduction, vocalization (in the case of animal kinds), as well as functional roles and perhaps cultural norms in the case of conventional or social kinds. The distinction between positive and negative counter-instances ultimately amounts to a difference in psychological salience:

A positive counterinstance to Ks are F occurs when an instance of the kind K has a concrete alternative property, that is, when it has a *positive alternative* to the property F, while negative counter-instances occur when an instance simply fails to be F. Whether a counterinstance counts as positive or negative is highly dependent on the property being predicated. (Leslie 2007: 66)

If a generic predicates a property with salient *positive* counter-instances, then the generic is likely to be judged false. If the counter-instances to the property predicated are negative, then the generic is likely to be judged true. For instance, there is a positive alternative to being *female*; one could be *male*. On the other hand, it is not clear that there is any positive alternative to *laying eggs*, at least for ducks. Thus, 'ducks lay eggs' is judged true because laying eggs lies along a characteristic reproduction dimension for ducks, and there are no positive counter-instances. Conversely, 'ducks are female' is judged false because there is a positive alternative property to being female.³

Generics of Type B, on the other hand, are intuitively true in virtue of ascribing 'striking properties' to kinds. Striking properties are harmful, frightening, or dangerous properties which we would be wise to avoid. Examples include 'Mosquitoes carry West Nile' and 'sharks attack swimmers.' These generics are frequently judged to be true even when the objective prevalence of the property among the kind is very low, since the cognitive cost of treating these generics as true is far lower than the possible risk of treating them as false. Consequently, people tend to overestimate the prevalence of a striking property among a relevant kind, and these

³However, see Sterken (2015) for counterexamples to and arguments against the claim that all true Type A generics have negative counter-instances.

sentences tend to receive a weighted role in inference (Khemlani, Leslie, and Glucksberg 2009). Moreover, the knee-jerk overgeneralizations which Type B generics prompt seem to play a role in the formation of prejudices and stereotypes about social kinds (Leslie forthcoming; Rhodes, Leslie, and Tworek 2012).

Leslie takes the behavior of Type A and Type B generics to be strong evidence for two claims. First, our facility with these types shows that generics express ‘cognitively fundamental generalizations’ about kinds. The seemingly disordered character of our semantic intuitions about tricky generics shows that this fundamental generalization mechanism is rife with biases which affect the space of kinds and properties over which we generalize. Second, the semantics of the GEN⁴ operator must be relatively simple, since even young children can grasp the semantic contribution made by the generic operator, even though it can be attached to all sorts of different sentences. A uniform metalanguage interpretation of GEN thus cannot account for the truth-conditional diversity that Type A and Type B generics display. Accordingly, Leslie opts to simply disquote GEN in the semantics, offering no interpretation of GEN in the metalanguage:

‘Tigers are striped’ is true iff [GEN] Tigers are striped.

In order to account for the behavior of Type A and Type B generics, Leslie (2007, 2008) offers a ‘weak semantics’ which specifies the set of conditions which any generic sentence must meet if it is to be judged true:

A generic K’s are F is true iff:

⁴According to the dominant trend in semantics, the logical form of generic sentences is very close to the logical form of sentences which include adverbs of quantification such as ‘usually,’ ‘generally,’ etc. (Lewis 1975; Krifka et al. 1995; Leslie 2007.) In particular, sentences of this sort have a *tripartite structure* which includes a variable-binding Operator of some sort, a Restrictor, and a Matrix:

Operator_x, ... ,z [Restrictor_x ... z] [Matrix_x ... z]

The Restrictor specifies the objects to which bound variables are assigned. The Matrix specifies the property or properties which are assigned to the objects in the Restrictor. The job of the Operator is to bind the variables in the Restrictor and Matrix and to relate them to one another. In the case of sentences like ‘All ravens are black’ or ‘Some students are responsible,’ the Operator will be a universal or existential quantifier. In the case of generics, the Operator will be a variable-binding Operator GEN. Bracketing the question of what the semantic contribution of GEN is, the logical form of ‘Tigers are striped’ will look like this:

(1b) GEN [Tiger(x)] [Striped (x)]

- (1) The counter-instances to the claim are *negative*. That is, there is no noteworthy or salient property G such that K's which are not F are G.
- (2) If F lies along a 'characteristic dimension' for K's, then some K's are F. If K is an artifact or social kind, then F is the 'function or purpose' of kind K.
- (3) If F is a striking property, then some K's are F, and all other K's are disposed to be F. If F is not a striking property, [and if 'K's are F' is not a minority Type A generic,] then almost all K's are F. Leslie is adamant that these clauses do not deliver *semantic* truth conditions for generics. Rather, they are 'worldly truthmakers' which specify how the world must be in order for the generic to be true.

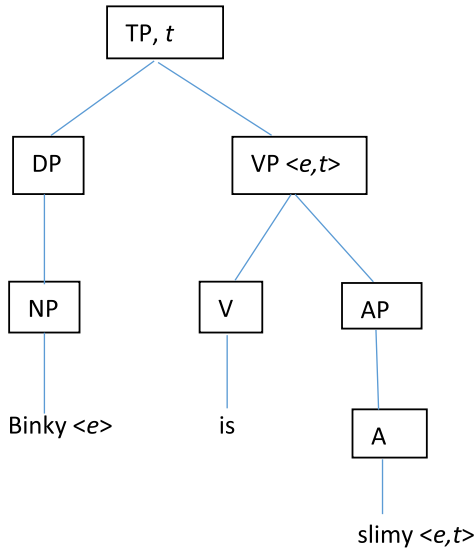
For Leslie, the distinction between semantic truth and worldly truthmakers is of the utmost importance. For one thing, she takes these truthmakers to be far more complex than the purely semantic clauses for generics. If the semantics of generics were as complex as these truthmakers, we would not expect children to display competence with generic interpretation and generic truth value judgments until much later in their development. More importantly, the distinction between the semantic question of what sentence *X* means and the metaphysical question of what the *world* must be like if sentence *X* is to be true seems to be rather fundamental.

For example, take the example of what it is to be *slimy*. Semantic questions would include: what does 'slimy' mean? What kinds of objects are in 'slimy's extension? Metaphysical questions would include: what facts about an object make it *slimy*? Is *sliminess* a property of the world external to our senses? Or is it something that only exists if there are creatures which can experience *sliminess*? Now consider the following sentence:

- (10) Binky is slimy.

In providing a semantics for (10), we can punt on the metaphysical questions of whether *sliminess* is 'out there' in the world, or somehow a 'product' of our experience, or whatever. All we're interested in is the sentence's compositional structure and its truth conditions. And these semantic features can be provided rather easily, using the standard resources of generative linguistics and first-order logic:

Compositional Structure of (10):



Truth conditions (10) Slimy(Binky)

On the other hand, if we are going to specify the conditions which must be in place if (10) is going to be true, we need to choose a specific metaphysical view of *sliminess*. For the sake of argument, let's say that in order for an object to be slimy, that object must be *experienced* as slimy by a normal observer in normal perceptual conditions. Call this 'the dispositional theory of sliminess.' Importantly, this theory says nothing about what 'slimy' *means*, but only says something about what it is for an object to actually *be* slimy. Unlike our analysis above, it doesn't give us the truth-conditions of (10), nor its compositional structure or logical form.

What is important for our purposes is that *even if* the dispositional theory of sliminess is correct, that would explain nothing about how we can be competent with the sentence 'Binky is slimy.' We can very well understand what the sentence says (once we fix a referent for 'Binky') without having any metaphysical theory of sliminess. If the dispositional theory of sliminess *did* deliver a semantics for 'slimy,' the theory would be part of what speakers would have to know in order to know that (10) is true. This is already quite implausible, since competent users of the term 'slimy' don't have to have any particular convictions about the metaphysics of sliminess. But to make matters worse, *even if* the metaphysical theory were to enter into the semantics, the logical

form – that is, the structure that feeds into compositional semantic interpretation – of ‘Binky is slimy’ would be:

$$\text{GEN } x, y \text{ [Standard Observer } (x) \wedge \text{Standard Condition}(y)] \\ \text{[Experiences As_In_}(x, \text{Binky, slimy, } y)]$$

The complexity of the clause makes it implausible that a child would learn how to interpret sentences structured this way before learning the much more simply structured clauses for quantifiers. More importantly, even if this were plausible, and there is a compelling argument for incorporating standard conditions and observers into the truth-conditions for (10), the formulation above delivers the wrong logical form for ‘Binky is slimy’. The GEN operator is needed because the dispositional theory dictates that Binky is slimy iff Binky is experienced as slimy across a *range* of conditions including more than one standard observer and more than one standard condition. But this is already problematic. ‘Binky is slimy’ is a statement of a particular fact about a particular individual, and as such, is patently non-generic. Leslie thus concludes that a semantic theory that takes us this far away from an empirically adequate account of the logical form and compositional structure of generics ought to be abandoned. The question I will pursue in the next few sections is whether this demarcation between semantics and metaphysics holds up in the case of Type C generics.

2. Introducing type C generics

All Type C generics have two interesting properties. First, Type C generics predicate what I call *socially perspectival predicates* (SPPs) of social kinds. Second, they represent social kinds as *non-accidentally* linked to the properties described by SPPs. We are here interested in a subset of Type C generics, namely those concerning racial kinds. Generics such as ‘Blacks are criminal/violent/thuggish’ and ‘Muslims are terrorists’ encourage *racialized* interpretations insofar as they lead interpreters to ascribe racial or ethnic characteristics to *criminality, violence, terrorism*, etc.

It is worth dwelling on this second feature of Type C generics, since the concept of racialization, native to sociology and political science, has received precious little attention within the semantics literature proper. ‘Racialization’ describes the process by which

racial or ethnic characteristics are ascribed to people, institutions, practices, and relationships within a social structure, frequently independently of the self-characterizations of the people, practices, etc., in question (Omi and Winant 2013).⁵ Consequently, crime is racialized to the extent that Blacks and other non-Whites are perceived to be especially disposed toward criminality,⁶ and terrorism is racialized (or at least indexed toward a specific religion) to the extent that Muslims (and especially Arab Muslims) are perceived to be paradigmatic examples of ‘terrorists.’

For the purposes of this paper, I largely take for granted that such interpretations are readily accessible. Other examples of racialized predicates include Ronald Reagan’s infamous rhetoric of the harmful self-entitlement wielded by ‘welfare queens,’ a term designed to elicit anti-Black and anti-welfare sentiments among White voters.⁷ More recently, writers in and outside academia have focused on the possibility that ‘thug’ is a racialized code for an intimidating, aggressive Black male.⁸ In my view, part of the job of a perspectival semantics is to explain how these SPPs receive these racialized interpretations.

The rest of the paper runs as follows. First, I introduce the central motivations of socially perspectival semantics and introduce the

⁵While it is not one of his words, perhaps the most wide-ranging and expansive discussion of the mechanisms and social effects of racialization is that of Fanon (2008).

⁶A wide variety of disciplines have produced excellent work on the causes and effects of the racialization of crime at institutional, interpersonal, and individual levels. Relevant facts include: an arbitrary Black person is over 7 times as likely to go to prison as an arbitrary white person (Alexander 2012), Blacks are about 2.7 times as likely to be subjected to vehicle investigative stops as whites (Epp, Maynard-Moody, and Haider-Markel 2014), Black teens are about 21 times as likely as Whites to be shot by police (Gabrielson, Jones, and Sagara 2014), over 50% of homicide offenders between 1980 and 2008 were Black (The U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics 2011), and both explicit and implicit psychological associations of Blackness with crime are shared by the majority of Americans (Watson, Jones, and Robinson-Saunders 1995; Correll et al. 2002; Nosek, Banaji, and Greenwald 2002, 2007).

⁷As it turns out, such appeals were effective. See Haney-Lopez (2014) for a book-length discussion of the role of these coded ‘dog whistles’ in American politics. One set of studies (Gilliam 1999) hypothesized that the term ‘welfare queen’ triggered a racialized cognitive and cultural script which held that Blacks (and especially Black women) are indolent and undeserving. The studies found that exposure to the ‘welfare queen script’ in visual media increased anti-Black bias, support for anti-welfare policies, and support for traditional gender roles.

⁸See Anderson (‘Notorious Thugs.’) Again, this insight is not confined to the academy. The popular sports news site *Deadspin* drew explicit attention to the way in which ‘thug’ conventionally tokens ‘Black’ after African-American professional football player and Stanford graduate Richard Sherman was widely accused of being a ‘thug’ after a boisterous post-game interview: ‘The word “thug” has been used so many times by the same sort of people about the same sort of thing that it’s no longer even accurate to call it code – it’s really more of a shorthand. It means a Black guy who makes White folks a little more uncomfortable than they prefer’ (Wagner 2014).

notion of a socially perspectival predicate (SPP.) I then show that Leslie's weak semantics for generics and its attendant distinction between semantic and metaphysical truth can explain neither (1) how SPPs (and a fortiori Type C generics) acquire their truth conditional semantic features nor (2) the underlying semantic structures that make default racialized interpretations of these generics plausible. Finally, I provide a sketch of an account that explains both of these aspects in the final section of the paper.

2.1. *Socially perspectival predicates*

According to received wisdom, some of our concepts 'carve nature at its joints.' Mountains, gold, water, and other mind-independent natural kinds are the joints, and the terms or concepts for these joints are the knives. We may bring the *concepts* of 'mountain' and 'water' to the world, but those concepts are answerable to how mountains and water actually are. The world's joints are metaphysically *perspective-independent* entities, and they are criterial for the correctness and applicability of terms/concepts which are *non-perspectival*.

Yet other concepts reflect the perspective of the categorizers, and indeed require such perspectives for their intelligibility. This is most obvious in the case of concepts for conventionally demarcated social kinds such as 'foreigner.'⁹ Unlike the set of mountains, which always consists of all and only those things which are mountains, the set of foreigners can only be assessed *from* a geopolitical perspective. I can even 'experience' the perspectival nature of 'foreigner' by going to a foreign country and realizing that *I* count as the foreigner!

'Foreigner' is an example of what I call a *socially perspectival predicate* (SPP.) We can characterize SPPs in terms of four interrelated conditions:

Socially perspectival predicates (SPP): Predicate P is socially perspectival just in case

- (1) Metaphysical condition: the property or object that P picks out is *perspective-dependent*.
- (2) Semantic condition: The extension of P is determined by a contextual parameter for a *relevant perspective* (RP).

⁹I borrow the example of 'foreigner' from Price 2007.

- (3) Pragmatic condition: patterned socially sanctioned uses of P *institute* the property or object which P picks out, and facts about these properties constrain and direct these patterned uses.
- (4) Epistemic condition: there is no social-practice-independent or non-conventional *epistemic* standard which can be used to determine whether a property or object X falls into the extension or anti-extension of P. The metaphysical condition entails that the properties described by SPPs do not exist independently of the perspectives of the creatures who apply such predicates. Independently of labeling people as ‘thugs,’ ‘terrorists,’ and ‘criminals,’ there would be no thugs, terrorists, or criminals – though there would be mountains, insects, quarks, and shapes, even if these objects were never labeled as such by a linguistic community.¹⁰

The semantic condition hypothesizes that perspectival terms like ‘foreigner,’ ‘Black,’ ‘criminal,’ and others contain in their semantics a contextual parameter that needs ‘filling’ by a *relevant perspective* (RP), in the same way that predicates of personal taste such as ‘tasty’ require an experiencer parameter in order to model their contribution to truth conditional content (e.g. Glanzberg 2007).

(11) [[tasty]] = *tasty-to-experiencer E*

(12) [[foreigner]] = *foreigner-for-relevant perspective RP*

The pragmatic condition suggests that the dominant patterned uses of SPPs are capable of instituting or creating social kinds. My account is officially neutral on the mechanisms that permit this institution (it can, for instance, accommodate various proposals by theorists of social construction), but the basic idea is that there is a feedback loop between practices of labeling or describing bits of the world and facts about the world external to such practices. In such conditions, labeling an object with P creates constraints on

¹⁰I take it that the existence of metaphysically perspective-dependent properties depend on the making of *actual* judgments about these properties, but I do not mean to foreclose the possibility that some such properties depend on there being *dispositions* to detect these properties, even if no perspective actually does detect them.

what kinds of things belong in the concept's extension, and facts about these objects mutually constrain the proprieties of use for P.

Finally, the epistemic condition entails that the procedures for determining whether or not an object fits into the extension of an SPP are set by the judgments, conventions, and practices of a linguistic community, not by any convention-independent standard. In order to find out what sorts of people in a given society counted as *criminal*, *thuggish*, or a *terrorist*, we would have to make concrete reference to social conventions such as the criminal code of that society, norms of etiquette and proper conduct, and the overarching political context. If we pair this epistemic thesis with the metaphysical, semantic, and pragmatic conditions, we get the result that a community's *judgments* and *practices* are authoritative over what kinds of objects actually fit into the extension of a given SPP. As long as the community has fixed an interpretation of what it means to be a 'criminal,' it cannot be wrong about who counts as *criminal*.

While there is frequently widespread agreement in a given social order about what sorts of things fall under a particular SPP and which do not, extensions can also be more-or-less locally contested. Contestations arise when members of a collective disagree on certain aspects of the publicly shared criteria for categorizing an object or kind with the predicate. Consider terms for human races, such as 'Black' or 'White.' Such terms are applied to different people according to a few loose criteria, including somatic characteristics (skin color, hair texture, etc.), prior ancestry, culturally specific 'Black' vs. 'White' experiences, personal racial identification, etc. (Mills 1998) More often than not, these criteria conflict when we try to categorize a given individual as White or Black (even assuming that we are only trying to categorize people of roughly European and/or African stock.) Borderline cases such as 'mixed-race' individuals and individuals capable of 'passing' as a member of another race show that these loose criteria deliver conflicting results depending on which aspects of the theory we regard as most relevant for any given case. In such cases, we might defer to a *meta-conventional* epistemic standard. For instance, one widely endorsed theory of racial membership – the 'one drop of blood rule,' according to which a certain percentage of 'Black' ancestry

is sufficient to make one 'Black' – frequently settles debate about whether a person is White or Black in contexts in which that meta-convention is honored.¹¹

As far as I can tell, the category of SPPs does not map neatly onto other linguistic categories that philosophers and linguists have studied.¹² However, SPPs seem to have some affinity with *predicates of personal taste* such as 'tasty,' 'delicious,' 'spicy,' 'disgusting,' 'fun,' and 'beautiful.'¹³ The properties these predicates pick out are metaphysically perspective-dependent: after all, nothing is tasty, fun, or beautiful unless there is some experienter *to whom* things are tasty, fun, or beautiful. Moreover, judgments about tastiness are authoritative over what kinds of things *are* tasty, at least to some extent. If you sincerely say 'cheeseburgers are tasty,' you can't be wrong about the fact that cheeseburgers are tasty *to you* – in this sense, your judgments about 'tasty' exhaust the scope of what kinds of things *are* tasty (to you.)

However, here the resemblances between personal taste and social perspectives fade. First, what you personally think is tasty or fun has no bearing on whether others in the community share your tastes: what seems tasty or fun to you might seem disgusting or deviant to your friend, but after all, *de gustibus non disputandum*. Second, it doesn't seem that the epistemic standards which are used to determine whether any given object is *tasty* are conventional, much

¹¹One might think that the one-drop-of-blood rule represents a convention-independent epistemic standard for 'Black' and 'White,' since what is being measured are non-perspective-dependent genetic facts of the matter. This is true, but it misses the point. The one-drop rule is a convention-independent epistemic standard which is invoked to justify a very much convention-dependent standard with a specific function: to subjugate those identified as 'Blacks' and to privilege those identified as 'Whites.' It is simply not the case that the one-drop rule identifies two different racial groups independently of the activities of humans and communities which draw racial distinctions for specific purposes.

For instance, a recent study (Bryc et al. 2015) found that about 13% of Louisiana residents who identify as 'White' have at least 1 'Black' ancestor over the last 11 generations (or 1% African ancestry.) The findings were hotly debated; did this make those white folks 'Black?' Meanwhile, there was very little discussion about a related topic: the average African-American has 24% European ancestry, or at least one 'White' ancestor over the last 3 or 4 generations. But no one seemed to wonder whether this made the average African-American 'White.' This would seem to confirm the existence of an implicitly agreed-upon conventional standard: one Black ancestor is sufficient to make you Black, but no amount of White ancestors is guaranteed to make you White.

¹²One obvious precursor is Gallie's (1955–1956) notion of an 'essentially contested concept.' The present account goes beyond Gallie's insofar as it attempts to diagnose the semantic, epistemic, and pragmatic mechanisms that lead to such contestation in the first place. Moreover, the category of SPPs arguably contains a wider swath of predicates than does Gallie's conception.

¹³This isn't to suggest that the category of 'predicates of personal taste' is any better defined than SPPs. The goal in each case is to take paradigm examples and see how far the linguistic data take us.

less social, in nature. Generally one knows immediately whether one is having an experience as of tasty or of something beautiful. One's own sensory apparatus, dispositions, and judgments are sufficient for determining whether a given *x* is tasty-to-one.

Even so, I take it that there are *uses* of so-called predicates of personal taste which more closely mirror the characteristics of SPPs.¹⁴ These uses do not report on what an object or experience is like to a particular experiencer, but seem to make a claim about how things are in the real world, independently of the perspective of any *particular* experiencer, but dependent on the perspective of possible *generic* experiencers in the linguistic community. If 'cool' can be used in this latter sense, I would be saying something false if I utter (13) in a community in which Beyoncé is widely and generically regarded as cool:

(13) Beyoncé is not cool.

On the other hand, (13) could be true in a context in which I am understood to be reporting on my own opinions about Beyoncé. This ostensibly dual function of 'cool' seems to give us reason to think that the first pair of sentences below describes a disagreement, while the second pair does not. Assuming the speaker is being truthful, a retraction or clarification of A's position following the first pair of sentences seems entirely natural, but not the second pair.

A: Beyoncé is not cool [generically].
 B: What are you talking about? Of course she is!
 A: Ok, well she's not cool *to me*. I have different tastes.

A: Beyoncé is not cool [to me].
 B: What are you talking about? Of course she is!
 A: ?? You're right. I lied before.

In the first set, the speaker backs down from the strong claim that Beyoncé 'lacks coolness,' so to speak, in order to report on her personal dispositions about what she finds cool. In the second set, the speaker begins with a report on her dispositions, and her interlocutor either (1) misunderstands that the speaker is using 'cool' in this private way and challenges her, (2) reports on his own dispositions regarding coolness, thus failing to disagree with the speaker, or (3) understands that the use is private, but maintains

¹⁴Strictly speaking, I am as yet agnostic on whether this is actually a difference in the use of a term, or whether this is evidence for there being more than one lexical entry for some predicates of personal taste.

that the speaker misunderstands what the proper extension of ‘cool’ is. In all three cases, the speaker’s retraction is infelicitous.¹⁵

This suggests that determining what (possibly generic) sets of judges regard, treat, or believe to be *cool* is vital to determining what ‘cool’ means in context. This is a point which is generally overlooked in the generics literature. For instance, Krifka et al. (1995) claim that semantic theories must be able to distinguish between what generics say about ‘the world’ and what they say about ‘cultural norms’ in order to be viable:

For instance, suppose it is the norm in some culture to assume that snakes are slimy. Even in that culture, the sentence *snakes are slimy* is a false sentence – although believed to be true by most members of the culture – since snakes, those real-world objects, are in fact not slimy. That is, generics are construed about making claims about the world, rather than what is considered a cultural norm. (Krifka et al. 1995, 49)

Here the assumption is that the purpose of generics is to describe the world as it is independently of our beliefs, judgments, and behaviors. While this is true of a wide variety of generics, generics which contain SPPs challenge this assumption. Replace ‘snakes are slimy’ in the above paragraph with ‘drugs are cool’ and see how it stands – ‘the sentence *drugs are cool* is a false sentence – although believed to be true by most members of the culture – since drugs, those real-world objects, are in fact not cool!’ The fact that most members of a culture believe that snakes are slimy is neither a necessary nor sufficient condition for snakes being slimy, but the fact that most members of a culture believe that drugs are cool is necessary and possibly sufficient for it being the case that drugs are cool, at least in the generic sense outlined above.¹⁶

However, even this doesn’t go quite far enough. ‘Drugs are cool’ isn’t true because *most* people think it’s true or because they want it to be true; rather, it’s true because some relevant *authority* thinks

¹⁵In the third case, perhaps B’s utterance could lead to a negotiation over which perspective on ‘cool’ is the most appropriate one, and perhaps A might eventually retract her statement after endorsing B’s conception of ‘cool.’ However, immediate retraction is odd, to say the least. (And even when it occurs, this situation can plausibly be read as feigned endorsement in order to curry favor with B rather than as a true change of semantic heart.) Thanks to an anonymous reviewer for pointing out this third interpretation.

¹⁶Another, perhaps more effective example: ‘Cows are food,’ is true if ‘food’ selects the dominant culinary-economic context in which cows are treated as things to be consumed, and the ability of ‘food’ to select this context is entirely dependent on whether or not cows are treated as food by some set of dominant practices. At the same time, ethical vegans and vegetarians can, with some justification, argue that the sentence is false at other, less dominant contexts.

or says so.¹⁷ In this sense, the extensions of SPPs can be decided in an *oligarchical* fashion. Krifka et al.'s requirement here is thus hampered all the more by the fact that it seems to rule out an analysis of virtually any generic that invokes a social or normative category. Part of what it is to belong in the extension of 'terrorist' or 'criminal' is to be *treated* like a terrorist or criminal in practice, and to be *believed* to be a criminal or terrorist by some relevant authority. That is, in applying the predicates 'criminal,' 'terrorist,' 'thuggish,' and so on, one needs to have a grasp of a certain relevant perspective from which terrorism, criminality, and thuggery are assessed – and it is part of the very idea of social perspectivalism that the relevant perspective will frequently not be one's own.

3. Leslie reconsidered

Type C generics pose two problems for Leslie's account. First, they challenge her distinction between semantic truth and metaphysical truth-makers by showing that in a wide range of cases, the things that must be the case in order for a generic to be true are sometimes the very conditions under which a predicate gets its semantic meaning. Those conditions frequently consist of *shared understandings* of what it is to *be* an X, and it is sometimes impossible to isolate these understandings from semantic questions about what 'X' means. Second, the default racialized readings of Type C generics cannot be accounted for within Leslie's weak semantics.

3.1. *Contra Leslie 1: perspectives and semantics*

Leslie's reasons for keeping truth conditions separate from worldly truthmakers are good ones. Semantic theorizing about what the sentences of a natural language mean should be constrained by facts about how human beings manage to pull off the trick of learning that language in the first place. And if children are able to display adult-like competence with using and interpreting generics even before they show competence with quantifiers, that is good evidence for the compositional semantic profile of generics being very simple. However, what we have seen is that even if this

¹⁷This 'externalist' commitment is important. We'll see an example of this oligarchical externalism when we turn to case of authority relations in baseball in the next subsection.

profile is simple, the worldly conditions that must hold in order for a given generic to be true are quite complex. In fact, they're too complex – the clauses that have to be built into an analysis of worldly truth makers might take us far away from the original compositional structure of a sentence (recall our 'Binky is slimy' example.)

However, it isn't clear that the metaphysical/semantic distinction can be made so quickly. Leslie's clauses are not only 'metaphysical truth-makers,' they are conditions under which a generic is *understood*. That is, while the clauses may have little to say about what the semantic features of generics are, they have a great deal to say about the epistemic question of how generics are *interpreted*. Which generics an interpreter judges true will be a function of how that interpreter perceives worldly conditions to be. If Leslie is right, and generics indeed are a source of evidence for how our primitive generalization mechanisms work, then it is plausible to think that these generalizations privilege different properties of kinds in judging certain generics true. For instance, the following generics are true, and the type of facts that seem to make them true (to an interpreter) differs in each case:

- (14) Turtles are long-lived. (*biological endowment*)
- (15) Dobermans have pointy ears. (*artifactual*)
- (16) Dobermans have floppy ears. (*natural/non-artifactual*)
- (17) Scotsmen wear kilts. (*social/cultural tradition*)
- (18) Prime numbers are divisible by one and themselves. (*axiomatic/definitional*)

Children and adults alike treat some types of worldly properties rather than others as more important to the evaluation of a given generic, depending on what specific kinds and ascribed properties are involved. And if these metaphysical properties have such a strong effect on whether a generic is true or false, then that seems to be sufficient reason to incorporate *metaphysical* theories of what makes a generic true into the semantics for the generic.¹⁸

To this extent, Leslie's worry about compositionality loses some of its bite. Even if it is implausible that children are tiny metaphysicians, fully equipped with dispositional theories of sliminess that does not show that children are completely ignorant of what kinds of worldly features make it the case that a generic is true. All we

¹⁸Gelman (2003) shows that children are far more likely to produce generics about 'natural' kinds than 'artifactual' kinds. She concludes that children rely on the deliverances of an implicit folk metaphysics in evaluating and producing generic claims. Prasada and Dillingham (2006) show that adults access different explanations of why a generic is true depending on their folk conception of what ontological class objects and kinds belong to.

need in order to treat the metaphysics as reasonably continuous with the semantics is the premise that children rely on their knowledge of the world in interpreting and assessing generic claims,¹⁹ and that they frequently defer to other members of their linguistic community in assigning meanings to terms.

The question, then, is what kinds of properties are relevant to the interpretation of the Type C generics we are tracking? A great deal will hinge on what properties are relevant to the interpretation of various SPPs. In the remainder of the paper, I motivate the claim that in order to make sense of the interpretation of SPPs like 'terrorist,' 'criminal,' and 'thug,' we need to allow the semantic dimensions of SPPs to interact with facts about their function within particular sets of social practices.

3.2. *Proof of concept: baseball*

In order to show that SPPs have the type of lexical complexity I am supposing they have, we need to show two things. First: that there *are* predicates whose extensions are determined by social, conventional, practice-dependent standards, and that those standards yield *epistemically authoritative* interpretations of the extension of those predicates. Second: that in this process of extension-fixing, a metaphysically response-dependent *kind or object* is instituted or created.

Let's start simple. Imagine a game of baseball.²⁰ The game is governed by two kinds of rules, *constitutive* and *regulative*.²¹

¹⁹See Prinz (2002) for a thorough defense of the claim that background knowledge can figure into semantic interpretation in a way which still obeys the compositionality constraint. I must remain agnostic on what counts as metaphysics being 'reasonably continuous' with semantics for the moment, but if Gelman's and Prinz's acquisition stories are workable in the case of social kinds, we have good grounds to challenge Leslie's dichotomy.

²⁰For those unfortunate enough to lack familiarity with the glorious ordered complexity of the game of baseball, I can only point to the official baseball rules (MLB Publications 2015). However, for the purposes of the example I set up, you only need to know the following: a player, known as the 'pitcher' throws the ball to his opponent, the 'batter'. This throwing event is called a 'pitch'. Every pitch is either a 'ball' or a 'strike'. Roughly speaking, a 'strike' is a pitch that the batter swings at or a pitch which crosses home plate and is within the 'strike zone' – an invisible area extending roughly from the batter's knees to his chest. Conversely, a 'ball' is a pitch that the batter does not swing at *and* which is outside the strike zone. (Baseball fans will already note the inadequacy of this characterization, but it should do for our purposes.) The pitcher's objective is to prevent the batter from reaching base, and the batter's objective is to reach base. One way for the pitcher to win this duel is to get three strikes past the batter, and one way for the batter to win this duel is to force the pitcher to throw four balls.

²¹The basic contrast between these types of rule has been unpacked in different ways by Rawls (1955), Searle (1995), Haugeland (1998), and many others.

Constitutive rules for a game or practice are rules which somehow 'constitute' a given game or practice. The constitutive rules of baseball make it the case that baseball games can exist: these rules specify what counts as a ball or strike, under what conditions a team earns a run, and that if a batted ball is caught, then the batter is out, etc. Regulative rules, on the other hand, 'regulate' an already existing practice: they specify appropriate rules of conduct for the players. Thus a batter who strikes out is obliged to leave the field and sit in the dugout, and a runner is obliged to run the bases if he is to score a run for his team. Obviously, constitutive and regulative rules interact. A player who conducts herself appropriately also manages to follow the constitutive rules of the game.

The rules of baseball manage to confer semantic contents on linguistic performances within the game and *also* manage to constitute baseball objects. Deciding just *which* baseball objects they are is largely a matter of what perspectives we are prepared to treat




Figure 1. The Strike Zone. Courtesy of Major League Baseball.

as authoritative. Let's take the example of strikes. When is a pitch a strike? There are at least three natural interpretations, each of which has substantial effects on how the game is played.

Let X be a pitch from pitcher to catcher.

Let Condition S be met just in case X passes through the *strike zone* (Figure 1).

Let U be the home-plate umpire. U 's job is to respond to any X with one of two calls: 'ball!' or 'strike!' Any X that is not a ball is a strike, and vice versa.

STRIKE₁: X counts as a STRIKE₁ just in case X meets condition S , regardless of whether or not U calls 'strike!' U 's call is in a sense entirely superfluous.

STRIKE₂: X counts as a STRIKE₂ just in case condition S is satisfied or U calls 'strike!' U 's calling 'strike!' is sufficient for X to count as a STRIKE₂, but it is not necessary. That is, it is possible for an X which satisfies condition S to count as a STRIKE₂, even if U does not call 'strike!'

STRIKE₃: X counts as a STRIKE₃ just in case U calls 'strike!', regardless of whether condition S is met. U 's calling 'strike!' is necessary and sufficient for X to be a STRIKE₃.

Each idealized interpretation has different implications for the epistemic authority that U 's call bears.²² For instance, if what it is to be a *strike* is to be an X of type STRIKE₁, U can be wrong, even systematically wrong, about whether any given X 'really' is a strike. In fact, U 's call is entirely superfluous. Strictly speaking, U doesn't play any special role in conducting the game; the question of whether X is a strike could be resolved by appeal to a replay device that measured precisely whether or not X satisfied condition S . At the other extreme, if what it is to be a *strike* is to be an X of type STRIKE₃, then it is *impossible* for an umpire to be wrong about whether a given X is a strike. His utterance, 'strike!', functions as a performative which *makes* it the case that X is a strike, much in the same way that saying 'I now pronounce you married' in the right context makes it the case that two people are now married. At this extreme, it does not even seem as if U is still bound by the constitutive rules: rather, he legislates them as he goes along.

We need an intermediate position which does justice to the notion that the umpire's call is answerable to how things stand in

²²This does not mean that it is easy to tell in practice which norm is in play. For instance, some things that count as STRIKE₂'s will look indistinguishable from STRIKE₁'s.

the world (that is, with the nature of X's) and to the notion that the umpire's call is authoritative over what those X's actually are (i.e. balls or strikes.) STRIKE₂ is just such a position. According to STRIKE₂, an umpire can be wrong *or* right about whether any given X is a strike, depending on which perspective we treat as authoritative. Say that an umpire calls an X that does not satisfy condition S. From the point of view of the constitutive rules laid out by the official rulebook, U is clearly incorrect about whether X was *really* a strike. However, from the point of view of the regulative rules, which specify how players and umpires are supposed to conduct themselves during a baseball game, the umpire's call is authoritative. As it turns out, the regulative rules stipulate that once an umpire makes a strike call, that judgment can be neither challenged nor overturned. The practice of strike-calling is thus functionally *oligarchical* in the sense noted above – every player on the field and every spectator could deny that the umpire's call was correct, but that wouldn't make the sentence 'that was a strike!' any less true.²³ In short, in baseball, the umpire's making a *mistake* in calling X a strike (according to the rulebook perspective) is completely compatible with X's *being* a strike (according to the perspective honored in practice.)

But wait, you might say, 'why assume that the *truth value* of 'strike!' can vary with which perspective is treated as relevant? I can see at least two ways in which this delivers the wrong result. First, you haven't given me an independent reason to treat the umpire's perspective as *semantically* relevant. Sure, it's relevant to the kinematics of the game – we can't proceed unless the umpire makes *some* call – but that's not sufficient for showing that the umpire's call is relevant to the *semantics* of 'strike.'

Second, and relatedly, I think I know what's going on here. You're assuming that authority over the kinematics of score gives the umpire a form of authority that is relevant to the semantics of 'strike,' but at best, umpires have an *epistemically* privileged position when it comes to calling strikes. Why not assume that the umpire is simply getting things wrong when he calls strikes which do not satisfy the rulebook notion of strike, full stop? That is, the umpire knows what 'strike' means – it just means what the constitutive rules say – but he is frequently mistaken about what kinds of things actually are strikes, much in the same way that I can misperceive a cow *as* a horse, even if I'm an expert on cow-identification. In fact, 'strike' calls *must* have

²³Of course, the perspectives of authoritative bodies that outrank the individual umpire – the Major League Umpire's Union, say, or Major League Baseball – might reverse this valence.

merely epistemic (and not semantic) significance. After all, umpires can only make so many mistakes until they are sanctioned by Major League Baseball. It's hard to see why such sanctions would be appropriate if the umpire speaks *truly* every time he calls a strike.

These are potent worries, and well taken. Indeed, I think that this alternative explanation is compatible with my own, and may be more acceptable for some semantic purposes (imagine the intractability of teaching a computer what a 'strike' is if my account delivers the One True Semantics!) However, I believe that the objections ultimately arise from an insufficient appreciation for how social practices among competent human language users can confer *functional* semantic properties on the terms of a language. What is a functional semantic property? Haslanger's (1995, 2005) illuminating vocabulary of 'manifest' and 'operative' concepts is helpful in understanding the basic idea. For Haslanger, a 'manifest' concept is explicit, public, and officially recognized. Applying manifest concepts is ultimately a casuistical problem. A candidate manifest concept-applier has to appeal to a more or less explicit codification in order to subsume a given part of the world under a concept. In the context of baseball, then, the manifest concept of a 'strike' is just what the official rule book says that it is. On the other hand, an 'operative' concept is more implicit; it is something of an 'unwritten norm' which is practiced despite its lack of official institutionalization. STRIKE₂ best captures the operative notion of a strike: while the umpire can be wrong about what a strike is from the perspective of the manifest interpretation, he cannot be wrong according to the operative interpretation – he makes the call, and the game goes on. In this case, it is the functional, operative notion that is privileged within the practice of baseball, even though it is clear what it is for any X to be a strike from the point of view of the manifest concept.²⁴

Second, the objection above does not so much argue for as *pre-suppose* that strikes are best thought of as perspective-independent

²⁴One might well wonder why the operative interpretation of strikehood is not STRIKE₃, according to which the 'strike' call functions as a performative which makes it the case that a given X counts as a strike, just as 'you are now married,' expressed in the right kind of circumstance, makes it the case that a given X and Y are married. The reason that this seems wrong is that this would entail that the constitutive rules play *no* role in constraining what kinds of X's can be strikes: an umpire can be systematically and utterly wrong according to the constitutive rules and still not be able to be sanctioned. However, this seems to represent a breakdown in baseball normal conditions. Major League Baseball seems to agree: they sanction their umpires if they are *systematically* out of step with what the rulebook specifies.

entities – that is, that there is a perspective-independent *fact* of the matter about whether or not the ball crossed some part of the plate, and thus we should draw a distinction between strike-seemings and actual strikes. However, this misses the point of the example: it's not that there are two notions of 'strike' and that one is a perspectival notion, it's that there are (at least) two notions of 'strike' and that *both* are perspectival. The argument is over *which* of these perspectives we should treat as relevant – the rulebook perspective or the umpire perspective? I have argued that the umpire perspective is the one which is semantically and epistemically relevant because of features of how baseball actually works in practice.

Finally, the fact that a perspective is 'in play' does not entail that everyone (or indeed anyone!) participating in the baseball game *consciously adopts* this perspective themselves, or knows that this is the perspective that is governing some aspect of game.²⁵ My account makes room for the possibility that even the umpire *himself* is not authoritative over which perspective is in play. In fact, the umpire, no doubt, thinks of himself as trafficking in the manifest concept – he takes himself to speak from the rulebook perspective. An umpire might even admit that he made a mistake after the game ends, suggesting that the umpire takes himself to be honoring the rulebook perspective, and expresses regret at failing to live up to it. However, on the externalist notion of perspectival determination we've set up, features of how baseball is played in non-ideal circumstances trump what baseball practitioners *take* to be the authoritative perspective. In this sense, the umpire perspective is authoritative *even when the umpire does not recognize it as his own perspective*. This possible mismatch between *acknowledged* speaker commitments on the one hand and *actual* speaker commitments on the other is especially important when we turn to an investigation of how racialized terms function within discursive practices – despite one's intentions to the contrary, it is possible to signal allegiance to a dominant perspective simply through the use of a word.²⁶

²⁵Thanks to an anonymous reviewer for drawing my attention to this possibility.

²⁶The possibility of such perspectival signaling has been taken quite seriously in recent years. Camp (2013), for instance, proposes an intriguing account of how such 'signaling' might work in the case of slurs.

3.3. Racial perspectives

What the objection gets right, I think, is that within the kinematics of a highly structured social practice like baseball, the choice of the perspective that is relevant to determining the semantic and metaphysics of 'strike'/*strike* is more or less arbitrary. My diagnosis is compatible with an alternative account that treats rulebook perspectives as authoritative. Arguably, that's because in baseball, *nothing much hinges* on what perspective we treat as relevant.²⁷ However, nailing down the perspective which determines the functional meanings of our racialized terms matters a great deal for racial ontology. Why? Because the functional properties of racialized talk very well *can* play a global role in organizing racial populations in a way that is highly consequential, and perhaps even dominating and unjust. Here the stakes are much higher, and the question of which perspective we treat as semantically and ontologically relevant is more important.

For instance, both the theorists and folk of oppressed groups have long known that dominant epistemic perspectives can dictate not only what sorts of entities count as *persons*, but also what different classes of *people are* in a functional sense. Consider even non-explicitly racialized terms like 'personhood,' 'citizenship,' and 'equality.' The famous opening words of the American Declaration of Independence, arguably the cornerstone of much of our manifest, official, publicly avowed thinking about human dignity, frame these categories as follows:

We hold these truths to be self-evident, that all men are created equal, that they are endowed by their Creator with certain unalienable Rights, that among these are Life, Liberty and the pursuit of Happiness. (Jefferson and Franklin 1776)

On the other hand, Blacks and other groups have long held privileged perspectives on what concepts such as *personhood*, *justice*, and *dignity* mean in their practiced, operative sense:

When White people say 'justice,' they mean 'just us.' (Black American folk saying, quoted in Mills 1997)

²⁷Or at least it doesn't matter in the sense of being highly relevant to ethical and social concerns *outside* the practice of baseball. The choice of baseball perspective of course does matter a great deal in the local world of high-stakes baseball contexts (no-hitters, perfect games, the World Series, etc.).

Race was far from irrelevant to [the Enlightenment conception of] personhood; skin color, hair, and facial features were used to categorize people and determine their moral standing ... that is the unacknowledged dark side of the Enlightenment ideal ... All persons are equal, but only White men are persons. (Mills 1998)

The lesson generalizes. When we reflect on the manifest uses of our concepts, we find that there are a lot of things we know about the social world, and a lot of things we agree upon. This comprises a rich shared stock of social knowledge. We know that anyone can be a *criminal* or *terrorist*, given the right circumstances. We know that every person in our society is entitled to life, liberty, and the pursuit of happiness. However, reflection on the operative dimensions of SPPs like ‘criminal,’ ‘terrorist,’ and ‘person’ tells another story. While anyone *can* be a criminal, Blacks are more criminal than Whites. While all persons are equal, Blacks are less equal than Whites.²⁸ While anyone can be a terrorist, the real terrorists are motivated by a fundamentalist interpretation of Islam.

Our stock of shared, explicit, ‘manifest’ knowledge has a murky underbelly. We implicitly rely on culturally distributed representations, beliefs, and assumptions that structure our interpretation of the social world and shape our social practices, sometimes resulting in profoundly unjust social arrangements. Many of these representations are racialized, and seep into terms in our language, thereby becoming part of the conventional meaning of certain terms – ‘thug,’ ‘welfare queen,’ ‘criminal,’ ‘terrorist.’²⁹

A natural question in the vicinity is *how* such terms come to mean what they do, and how we manage to have this implicit knowledge. This introduces a number of questions about the relationships between relevant perspectives and semantic values. What ensures that a given perspective will be the relevant one?

²⁸In fact, the operative and manifest notions of personhood came into sharp concordance during the ‘Three Fifths Compromise’ of 1787, which treated each slave as 3/5ths of a person in tallying the total population of a state in order to assign the appropriate number of representatives. Southern states thus received more representation than they would have had had slaves not been counted. The result was an overrepresentation of slaveholder interests in Congress until 1865.

²⁹It is important to note that not all racialized terms are racialized in the same way at all contexts, and that identical terms might have different RP values in different contexts. For instance, ‘thug’ is frequently used in ‘reappropriative’ contexts in order to specify and celebrate positive aspects of Black male identity. (See Anderson, [forthcoming](#)) I read such cases as instances of social negotiation over what perspective ought to be authoritative and honored. (Consider a White person hurling ‘thug’ as an insult at a Black man, and the latter taking it as a compliment.)

Once determined, how do relevant perspectives become authoritative over the meanings and uses of our terms? How broad is this authority, and what determines whether or not a given RP is at play within a discursive context? These are difficult and exciting 'metasemantic' questions. While I doubt that there is a single answer to these questions, and that each calls for case-by-case empirical work, I do think that our discussion so far has given us some idea of where we should look for answers.

Most centrally, our baseball example suggests that we ought to be attuned to the ways in which *power* and *authority* function within discursive practices. In practices like baseball, which perspectives are relevant and authoritative is determined partially by the rules of the game (e.g. the rules stipulate that the umpire is the authority) and partially by the way that the practice plays out in non-ideal circumstances (e.g. the relevant, authoritative perspective is determined by a combination of the umpire's perspective and the 'operative' perspective that is observed in practice by the game's participants.) Things are not so clear in more complicated discursive contexts, however, partially because there are seldom explicit authority-conferring rules which make clear what perspective is in play.

Frequently, determining a relevant perspective is neither a top-down nor bottom-up process. More often, perspectival determination is a 'loopy' process in which top-down authoritative mechanisms specify and reinforce the relevant perspective, and bottom-up deferential patterns become reflected in large-scale institutional structures. For instance, the relevant perspective that determined membership within the political category of *personhood* is not solely the result of the 'top-down' legislation that explicitly denied political and moral standing to non-Whites and non-males, nor is it solely the result of the 'bottom-up' individual prejudices and material interests of powerful, landowning Whites. Rather, each of these forces (and others) played a role in the *constitution* of a perspective from which the extension of 'personhood' could be determined. Moreover, the perspective that was so constituted was highly cross-contextual: Blacks were not only non-persons in most legal contexts, but also in most economic, social, and political contexts. Once established, the mechanisms that ensure that a certain perspective is treated as authoritative in practice can cut across a variety of domains. For instance, the mechanisms

that underlie a widespread deference to the perspective associated with ‘thug’ might include cognitive associational mechanisms which pair Blackness with violence or criminality, commitment to or acceptance of explicitly racialized political discourse, familiarity with negative stereotypes disseminated by popular media, and the institutional criminalization of Black (male) conduct.

These issues are deserving of much more discussion than I have the space to give them here. Assuming that a satisfying metasemantic and empirical story can be told about the various mechanisms that determine the nature and authority of relevant perspectives, we turn our attention to how the semantic mechanisms we have posited make racialized interpretations possible.

3.4. *Contra Leslie 2: racialized interpretations and weak semantics*

We’ve already seen that Leslie precludes the possibility of incorporating authoritative perspectives – an RP parameter – into the semantics for Type C generics and SPPs. The notion of racialized predicates gives Leslie’s weak semantics a test: if her account can successfully account for the interpretation of racialized SPPs, then it becomes less clear that SPPs require an RP parameter. However, I contend that her semantics fails this test for the same reasons as before: in treating semantic truth and metaphysical truth as isolated from one another, Leslie’s weak semantics cannot give an adequate account of either the meaning or metaphysical truth makers of Type C generics with racialized SPPs.

To see this, let’s revisit Leslie’s weak semantics for Type B generics, which is organized around the following ‘striking property clause.’³⁰

Striking Property Clause: If F is a striking property, then some K’s are F, and all other K’s are disposed to be F. If F is not a striking property, [and if ‘K’s are F’ is not a minority Type A generic,] then almost all K’s are F.

³⁰One might think that Leslie accounts for these readings via her ‘social kind’ clause: ‘if K is a social kind, then F is the *function or purpose* of the kind K.’ However, I think it implausible, to say the least, that *criminality* or *terrorism* are primarily conceived as (much less actually are) the function or purpose of blacks and Muslims. If such things were thought to be functions or purposes of these populations, for instance, it would be difficult to *fault* such kinds for performing those functions (although one would still be able to argue that those functions are bad.) In any case, there is evidence (Prasada and Dillingham 2006) that social kind generics are not normally given this teleological interpretation.

Type B generics thus admit of the following paraphrase:

- (19) Mosquitoes carry West Nile.
- (19a) Some mosquitoes carry West Nile, and all are disposed to carry West Nile.

Yet such paraphrases do not adequately capture the nature of the relationship that obtains between kinds and predicates in some Type C generics. Consider:

- (20) Blacks are criminal.
- (20a) Some Blacks are criminal, and all Blacks are disposed to be criminal.
- (21) Muslims are terrorists.
- (21a) Some Muslims are terrorists, and all Muslims are disposed to be terrorists.

If racialized interpretations are as accessible as I have suggested, these paraphrases are far too weak, regardless of whether they are supposed to be 'semantic' or 'metaphysical' clauses. Assume that they are metaphysical truth-makers. Thus the 'some' paraphrase is supposed to specify what the world must be like if an interpreter is to judge the generic true. Yet this would not explain why (20) and (21) seem to exert some inductive or cognitive 'pull' on many interpreters while (22) and (23) do not seem to exert such a pull.³¹

- (22) Some lawyers are criminal, and all lawyers are disposed to be criminal.
- (23) Some Swiss people are terrorists, and all Swiss people are disposed to be terrorists.

The cognitive 'pull' in question is that interpreters are likely to rely on representations of Blacks as criminal and Muslims as terrorists even despite few to no negative experiences with blacks or Muslims. The majority of Americans grossly overestimate the objective share of crimes committed by Blacks (Blow 2014), as well as the share of terrorist activity committed by Muslims (Ahmed 2015), and there is evidence that interpreters only need very scant evidence for Black criminality or Muslim terrorism in order to tacitly endorse a generalization to the effect that 'Blacks are criminal' or 'Muslims are terrorists' (Leslie *forthcoming*) That is, Blacks and Muslims are perceived to be *more disposed* toward criminality and terrorism than, say, lawyers and Swiss people.

However, the paraphrases do not explain why the bar for accepting 'Blacks are criminal' is so low, and why the bar for accepting 'lawyers are criminal' is so high. After all, the objective prevalence

³¹Leslie *forthcoming* has argued convincingly that generics like 'Blacks are criminal' and 'Muslims are terrorists' are structurally similar to Type B generics, at least in their psychological function (that is, these generics are closely linked to essentializing, stereotyping, and prejudicial attitudes.) So while I disagree with Leslie's semantic account in the case of Type C generics, I am happy to treat Type C and Type B generics as very closely related in their respective psychological functions.

of crime among lawyers is comparable (and perhaps greater) than the prevalence of crime among Blacks, but presumably crime among Blacks is given the greater weight. Of course, Leslie might argue that her metaphysical truthmakers simply specify minimal necessary conditions which must be met in order for the generic to be judged true, and that specifying any interpreter biases are not the job of her conditions. While I agree that these conditions are intended to be 'minimal' in this sense, I disagree that there *can* be any minimal specification of what the world must be like when it comes to the SPPs we're targeting, and especially when those predicates interact with racial kinds. That is, the question of who counts as a 'terrorist' cannot be delivered in any perspective-neutral fashion, and which perspective we choose will already come along with certain standards for who or what counts as one. Moreover, I take it that the facts about who counts as a terrorist (from some relevant perspective) *is* a feature of the world which is relevant for fixing the conditions under which a Type C generic is interpreted, and consequently should be part of Leslie's clauses. So, as I have said above, Leslie's demarcation is due to an inadequate division of labor between *what the truth conditions of a sentence are* and *what the world must be like in order for the sentence to be true*. 'Blacks are criminal' isn't judged true because 'some Blacks' are criminal; rather, it is judged true because its interpretation (and perceived truth-conditions) is sensitive to what it means to be a criminal given a set of ostensibly 'normal' racialized social conditions.

On the other hand, if (contrary to Leslie's intentions) these paraphrases give the *truth conditions* of generics, notice that the 'some' paraphrase remains true under virtually any substitution of a social kind; as long as one member of the kind has the property ascribed, the sentence is true. The paraphrases thus obscure the sense in which some generics receive default essentialized or racialized interpretations, thereby losing the sense in which being criminal or being a terrorist is perceived to be proper to certain kinds rather than others.

4. Doing better: a semantics for 'criminal'

So what kind of semantic machinery *do* we need to capture the racialized reading of Type C generics? In this final section, I

sketch the beginnings of an answer by focusing on the example of 'criminal'.³²

Luckily, the semantics literature already has a good deal to say about adjectives like *criminal*. *Criminal* belongs to the class of gradable adjectives, which includes words like *tall*, *short*, and *smart*.³³ Gradable adjectives can be paired with degree modifiers like *more*, *very*, *somewhat*, *kind of*, etc.:

- (24) Tim is very tall.
 (25) Alicia is kind of smart.
 (26) Murdering someone is more criminal than stealing.

According to one orthodox analysis (Glanzberg 2007; Kennedy 2007), gradable adjectives map the objects that they apply to onto *scales* relevant to the given gradable. So, for instance, 'tall' in (24), when modified by 'very,' maps Tim onto the 'upper end' of the scale for height, whereas 'smart' in (25), when modified by 'kind of' maps Alicia onto the more-or-less 'middle range' for intelligence.

Strictly speaking, a scale (S) is a triple consisting of three elements: *degrees* ($\delta_1 \dots \delta_n$), a *total ordering* of those degrees (O), and a *dimension* (D).

$$S = \langle \delta \dots \delta_n, O_{\delta_1 \dots \delta_n}, D \rangle$$

More precisely, then, the meaning of a gradable adjective is a function from individuals to *degrees* of the scale. Degrees can be thought of as the degree to which a given object is tall, smart, criminal, or whatever. The dimension provides 'a property according to which the degrees are ordered' (Glanzberg 2007). 'Tall' will thus be associated with a dimension that provides *height*, 'expensive' will be associated with a dimension that provides *cost*, etc. Finally, the dimension orders the degrees ($\delta_1 \dots \delta_n$) *totally*, meaning that relations between degrees on the scale respect antisymmetry, transitivity, and totality.³⁴

³²Obviously different treatments will be needed for different predicates. For instance, noun SPPs like 'terrorist' clearly won't be amenable to the approach to gradable adjectives I recommend here. However, the goal in each case will be the same: to show that the RP parameter makes a systematic contribution to the default interpretations of Type C generics and to their truth conditions. In other work, for instance, I am investigating the possibility that count nouns like 'thug' and 'terrorist' are *inherent generics* (Chierchia 1995) which contain an RP parameter in their semantic gloss.

³³I have encountered resistance to the idea that 'criminal' is truly a gradable adjective in addition to being a count noun. If you share these reservations, the term 'violent' will do just as well.

³⁴(a) Antisymmetry: If $\delta_1 \leq \delta_2$ and $\delta_2 \leq \delta_1$, then $\delta_1 = \delta_2$.

(b) Transitivity: If $\delta_1 \leq \delta_2$ and $\delta_2 \leq \delta_3$, then $\delta_1 \leq \delta_3$.

(c) Totality: Either $\delta_1 \leq \delta_2$ or $\delta_2 \leq \delta_1$.

The comparative form of gradable adjectives (e.g. ‘taller,’ ‘more criminal,’ ‘smarter,’ etc.) can thus be represented as a function which maps an individual object or kind onto some degree of tallness, criminality, or intelligence that is greater than the degree of tallness, criminality, or intelligence possessed by some other individual or kind:

- (27) (a) Tim is taller than Alicia.
 (b) The degree of height of Tim is tall is greater than the degree of height of Alicia.
 (c) $\text{height}(\text{Tim}) > \text{height}(\text{Alicia})$

In the absence of an explicit comparison class (i.e. in a sentence like ‘Tim is tall’), the meaning of the positive form of a gradable adjective is a function from individuals to some degree that *exceeds some contextually determined standard for height*, $s(\text{height})$:

- (28) (a) Tim is tall.
 (b) $\text{height}(\text{Tim}) > s(\text{height})$

What determines whether or not Tim exceeds the standard for height? According to this orthodox analysis, the conventional meaning of the adjective plays an important role in determining this. Roughly, according to the analysis in (28b):

Here s is a contextually determined function which picks out a contextually significant degree, based on the interpretation of a given gradable adjective ... [The] function s returns a contextually significant degree of tallness for input *tall*. The degree d is the value of s in context c . I shall refer to this degree value ... as the *standard* for that adjective and context (Glanzberg 2007: 9).

To be ‘criminal’ is to be located on a certain scale for criminality in a way that exceeds whatever the relevant standard for criminality is in a given context. Yet there are still several different ways ‘criminal’ can be interpreted, just as there are several ways ‘strike’ can be interpreted. Luckily, we can illuminate some of these interpretations if we revisit the hypothesis that SPPs contain an RP parameter which determines what objects fit into the extension or anti-extension of these terms.

If my hypothesis is on the right track, ‘criminal’ contributes a parameter which selects a relevant perspective (RP). As we saw, the reason that an RP belongs in the semantics is that SPPs describe perspective-dependent properties: in order for anything to count as a criminal, there needs to be some perspective from which a given person is judged to be criminal. In other words, one is only

a criminal *for* some code, social structure, or judicial system.³⁵ Moreover, as we saw in the case of 'cool,' RP can select collective or generic perspectives. Which sets of perspectives are actually selected will depend on features of the racialized social structure in which the term is current.³⁶ The meaning of 'criminal' in context thus receives the following analysis:

$$(29) \quad [[\text{criminal}]]^c = \text{degree-criminality-for-RP}$$

Once we introduce the RP parameter, we have arrived at one possible take on the truth conditions for 'Blacks are criminal':

- (30) (a) Blacks are criminal.⁶
 (b) The degree of criminality possessed by Blacks as a kind according to a relevant perspective is greater than the standard degree of criminality according to that perspective.
 (c) $\text{GEN criminal}_{\text{RP}}(\text{Blacks}) > s(\text{criminal}_{\text{RP}})$

Again, in providing these truth conditions, we have not yet specified the *particular* perspective which 'criminal' picks out. In the contexts we're interested in, those perspectives are selected by facts about racialized social structure, just as baseball perspectives are selected by facts about how baseball is conducted (see note 29.)

This semantics makes clear that RPs play a role both in determining the extension of SPPs and in determining the *standard* in virtue of which the properties they describe are associated with race. I have suggested that 'criminal' is scale invariant – the scale for criminality does not change depending on what social group 'criminal' is predicated of. In order to capture the racialized reading, I hypothesize that the degree value for 'criminal' required for 'Blacks are criminal' to be true is lower than the degree value required for 'lawyers are criminal' to be true. Blacks tend to be mapped to the upper end of the scale by default.

³⁵One source of linguistic evidence that an RP parameter belongs in the semantics is that 'criminal' and 'terrorist' can be complemented by *to/for/according to* adjuncts, as in 'Osama bin Laden was a terrorist according to everyone in the United States [but was a freedom-fighter according to others]' or 'Pot-smokers are criminal to Turkish police, but not to Dutch police.'

³⁶The question of what features these turn out to be is a question for what philosophers and linguists sometimes call 'metasemantics.' Figuring out *why* SPPs mean what they do introduces questions about the mechanisms by which 'criminal' has come to be symbolically connected with Blackness in America. The semantic approach I offer here is compatible with many such mechanisms.

Of course, this is not the only way to capture the racialized reading.³⁷ For instance, there could be semantic treatments which allow the scale associated with ‘criminal’ to vary. Words like ‘beautiful’ and ‘tasty’ can be associated with different scales. A painting might be beautiful according to a scale for *colorful* or according to a scale for *austerity*, and a beer can be tasty according to a scale for *darkness*, *refreshingness*, *bitterness*, etc. In the same way, it might be the case that sentences like ‘Blacks are criminal’ or ‘lawyers are criminal’ select different scales for criminality. For instance, perhaps ‘criminal’ selects *violent crime* or *drug crime* when it is predicated of blacks, and ‘tax fraud’ or ‘white collar crime’ when predicated of Whites.³⁸

5. Conclusion

This paper has made a case for the viability of a socially perspectival semantic framework for generics involving SPPs. I have provided two respects in which Leslie’s taxonomy of generics is unable to account for the interesting semantic features of this class, and motivated my own positive proposal about how psychological, social, and semantic mechanisms can interact in a way that delivers informative truth conditions for Type C generics. I hope to have shown that a socially perspectival semantics for generics can be extended to other semantic problems within philosophy and linguistics, and that semantics has important analytical tools to offer to sociology, political science, and various approaches to race and racism – and that cross-pollination among these disciplines is even more fruitful.

Acknowledgments

Thanks to Rachel Katherine Sterken, Bernhard Nickel, and Steven Gross. Thanks also to an anonymous reviewer for helpful comments.

³⁷One reason I prefer the scale-invariant option is that it seems to be able handle comparative generics of the form ‘Blacks are more criminal than Whites.’ If the scale were to vary (and, for example, ‘Whites’ were associated with a different form of crime than Blacks), then the truth conditions of the comparative construction would be uninformative: ‘Blacks are more criminal-for-Blacks than Whites.’ For this reason, I hope to maintain some ‘core meaning’ of criminal which allows us to make such comparisons.

³⁸Another scale-variant option would be that ‘criminal’ selects different *perspectives* depending on which group criminality is predicated of. ‘Blacks are criminal’ would thus select judges relevant to street crime, while other contexts would select judges relevant to tax fraud, white collar crime, etc. In this way, the scale would indirectly shift as a result of directly shifting the RP parameter.

Disclosure statement

No potential conflict of interest was reported by the author.

References

- Ahmed, Beenish. 2015. "Less than 2 Percent of Terrorist Attacks in the EU are Religiously Motivated." *ThinkProgress*. Accessed January 9. <http://thinkprogress.org/world/2015/01/08/3609796/islamist-terrorism-europe/>
- Alexander, Michelle. 2012. *The New Jim Crow: Mass Incarceration in the Age of Colorblindness*. New York: The New Press.
- Anderson, Luvell. *Forthcoming*. "Notorious Thugs." In *Black and Male: Critical Voices From Behind the Racial Veil*, edited by George Yancy.
- Asher, Nicholas, and Michael Morreau. 1995. "What Some Generic Sentences Mean." In *The Generic Book*, edited by Greg Carlson and Francis J. Pelletier, 300–339. Chicago, IL: Chicago University Press.
- Blow, Charles. 2014. "Crime, Bias, and Statistics." *New York Times*, September 7. <http://www.nytimes.com/2014/09/08/opinion/charles-blow-crime-bias-and-statistics.html>.
- Brandone, Amanda C., A. Cimpian, S. J. Leslie, and S. A. Gelman. 2012. "Do Lions Have Manes? For Children, Generics are About Kinds Rather Than Quantities." *Child Development* 83 (2): 423–433.
- Bryc, Katarzyna, E. Durand, J. M. Macpherson, D. Reich, and J. M. Mountain. January 8, 2015. "The Genetic Ancestry of African Americans, Latinos, and European Americans across the United States." *The American Journal of Human Genetics* 96: 37–53.
- Camp, Elisabeth. 2013. "Slurring Perspectives." *Analytic Philosophy* 54 (3): 330–349.
- Carlson, Greg, and F. J. Pelletier, eds. 1995. *The Generic Book*. Chicago, IL: Chicago University Press.
- Chierchia, Gennaro. 1995. "Individual-level Predicates as Inherent Generics." In *The Generic Book*, edited by Greg Carlson and Francis J. Pelletier, 176–223. Chicago, IL: Chicago University Press.
- Cohen, Ariel. 1999. *Think Generic!: The Meaning and Use of Generic Sentences*. Stanford: CSLI Publications, Center for the Study of Language and Information.
- Correll, Joshua, B. Park, C. M. Judd, and B. Wittenbrink. 2002. "The Police Officer's Dilemma: Using Ethnicity to Disambiguate Potentially Threatening Individuals." *Journal of Personality and Social Psychology* 83 (6): 1314–1329.
- Correll, Joshua, B. Park, C. M. Judd, and B. Wittenbrink. 2007. "The Influence of Stereotypes on Decisions to Shoot." *European Journal of Social Psychology* 37: 1102–1117.
- Declerck, Renaat. 1986. "The Manifold Interpretations of Generic Sentences." *Lingua* 68: 149–188.
- Epp, Charles R., Steven Maynard-Moody, and Donald P. Haider-Markel. 2014. *Pulled Over: How Police Stops Define Race and Citizenship*. Chicago, IL: Chicago University Press.

- Fanon, Frantz. 2008. *Black Skin, White Masks*. New York: Grove Press.
- Gabrielson, Ryan, R. G. Jones, and E. Sagara. 2014. "Deadly Force, in Black and White." *ProPublica*. Accessed October 10. <http://www.propublica.org/article/deadly-force-in-black-and-white>.
- Gallie, W. B. 1955–1956. "Essentially Contested Concepts." *Proceedings of the Aristotelian Society New Series*. Vol. 56, 167–198.
- Gelman, Susan. 2003. *The Essential Child*. Cambridge: MIT Press.
- Geurts, Bart. 1985. "Generics." *Journal of Semantics* 4: 247–255.
- Gilliam, Franklin. 1999. "The 'Welfare Queen' Experiment: How Viewers React to Images of African–American Mothers on Welfare." *Nieman Reports of the Nieman Foundation for Journalism at Harvard University*. Vol. 53, 2. <http://niemanreports.org/articles/the-welfare-queen-experiment/>
- Glanzberg, Michael. 2007. "Context, Content, and Relativism." *Philosophical Studies* 136: 1–29.
- Haney-Lopez, Ian. 2014. *Dog Whistle Politics*. New York: Oxford University Press.
- Haslanger, Sally. 1995. "Ontology and Social Construction." *Philosophical Topics* 23 (2): 95–125.
- Haslanger, Sally. November 2005. "What are We Talking About?: The Semantics and Politics of Social Kinds." *Hypatia*. 20 (4): 10–26.
- Haslanger, Sally. 2011. "Generics, Ideology, and Common Ground." *Feminist Metaphysics*, edited by Charlotte Witt, 179–207. Springer-Verlag.
- Haslanger, Sally. 2012. *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.
- Haugeland, John. 1998. "Truth and Rule-following." In *Having Thought*. Cambridge: Harvard University Press.
- Halberda, Justin, and Lisa Feigenson. 2008. "Developmental change in the acuity of the 'number sense': The approximate number system in 3-, 4-, 5-, and 6-year-olds and adults." *Developmental Psychology* 44 (5): 1457–1465.
- Jefferson, Thomas, and Benjamin Franklin. 1776. Preamble to the American Declaration of Independence.
- Kennedy, Chris. 2007. "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30: 1.
- Khemlani, S., S. J. Leslie, and S. Glucksberg. 2009. "Generics, Prevalence, and Default Inferences." Proceedings of the 31st Annual Cognitive Science Society. Amsterdam: Cognitive Science Society.
- Krifka, Manfred, F. J. Pelletier, G. Carlson, A. ter Meulen, G. Link, and G. Chierchia. 1995. "Genericity: An Introduction." In *The Generic Book*, edited by Greg Carlson and Francis J. Pelletier, 1–124. Chicago, IL: University of Chicago Press.
- Leslie, S. J. 2007. "Generics and the Structure of the Mind." *Philosophical Perspectives* 21 (1): 375–403.
- Leslie, S. J. 2008. "Generics: Cognition and Acquisition." *Philosophical Review* 117 (1): 1–47.
- Leslie, S. J., and S. A. Gelman. 2012. "Quantified Statements are Recalled as Generics: Evidence from Preschool Children and Adults." *Cognitive Psychology* 64: 186–214.

- Leslie, S. J. 2014. "Carving Up the Social World with Generics." *Oxford Studies in Experimental Philosophy* 1: 208–231.
- Leslie, S. J. **Forthcoming**. "The Original Sin of Cognition: Fear, Prejudice, and Generalization." *Journal of Philosophy*.
- Lewis, David. 1975. "Adverbs of Quantification." In *Formal Semantics of Natural Language*, edited by E. L. Keenan, 178–188. Cambridge University Press.
- McConnell-Ginet, Sally. 2003. "What's in a name?: Social Labelling and Gender Practices." In *The Handbook of Language and Gender*, edited J. Holmes and M. Meyerhoff, 69–97. Oxford: Blackwell.
- Mills, Charles. 1997. *The Racial Contract*. Ithaca: Cornell University Press.
- Mills, Charles. 1998. *Blackness Visible: Essays on Philosophy and Race*. Ithaca: Cornell University Press.
- Nickel, Bernhard. 2009. "Generics and the Ways of Normality." *Linguistics and Philosophy* 31: 629–648.
- Nosek, B. A., M. R. Banaji, and A. G. Greenwald. 2002. "Harvesting Implicit Group Attitudes and Beliefs from a Demonstration Web Site." *Group Dynamics: Theory, Research, and Practice* 6 (1): 101–115.
- Office of the Commissioner of Baseball. 2015. *Official Baseball Rules*. New York: Major League Baseball Publications.
- Omi, Michael, and Howard Winant. 2013. *Racial Formation in the United States*. New York: Routledge.
- Prasada, S., and E. Dillingham. February 2006. "Principled and Statistical Connections in Common Sense Conception." *Cognition*. 99 (1): 73–112.
- Price, Huw. 2007. "Causal Perspectivalism." In *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited*, edited by Huw Price and Richard Corry, 250–292. Oxford University Press.
- Prinz, Jesse. 2002. *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge: MIT Press.
- Rawls, John. 1955. "Two Concepts of Rules." *The Philosophical Review* 64 (1): 3–32.
- Rhodes, M., S. J. Leslie, and C. Tworek. 2012. "Cultural Transmission of Social Essentialism." *Proceedings of the National Academy of Sciences* 109 (34): 13526–13531.
- Searle, John. 1995. *The Construction of Social Reality*. New York: Free Press.
- Sterken, Rachel Katherine. 2015. "Leslie on Generics." *Philosophical Studies* 172 (9): 2493–2512.
- The U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Statistics. 2011. "Homicide Trends in the United States, 1980–2008." Accessed November. <http://www.bjs.gov/content/pub/pdf/htus8008.pdf>
- Wagner, Kyle. 2014. "The Word 'Thug': Regressing." *Deadspin*. Accessed January 21. <http://regressing.deadspin.com/the-word-thug-was-uttered-625-times-on-tv-yesterday-1506098319>.
- Watson, Betty, D. Jones, and P. Robinson-Saunders. 1995. "Drug Use and African Americans: Myth versus Reality." *Journal of Alcohol and Drug Education* 40 (2): 19–39.