

Special Section

Open Access

Theodore Bach

Social Categories are Natural Kinds, not Objective Types (and Why it Matters Politically)

DOI 10.1515/jso-2015-0039

Abstract: There is growing support for the view that social categories like *men* and *women* refer to “objective types.” An objective type is a similarity class for which the axis of similarity is an objective rather than nominal or fictional property. Such types are independently real and causally relevant, yet their unity does not derive from an essential property. Given this tandem of features, it is not surprising why empirically-minded researchers interested in fighting oppression and marginalization have found this ontological category so attractive: objective types have the ontological credentials to secure the reality (and thus political representation) of social categories, and yet they do not impose exclusionary essences that also naturalize and legitimize social inequalities. This essay argues that, from the perspective of these political goals of fighting oppression and marginalization, the category of objective types is in fact a Trojan horse; it looks like a gift, but it ends up creating trouble. I argue that objective type classifications often lack empirical adequacy, and as a result they lack political adequacy. I also provide, and in reference to the normative goals described above, several arguments for preferring a social ontology of natural kinds with historical essences.

Keywords: Natural kind; Objective type; Gender; Historical lineage; Empirical adequacy; Semantic externalism; Race; Conceptual engineering.

Theodore Bach, Associate Professor of Philosophy, Bowling Green State University, Firelands College, 1 University Dr, Huron Ohio 44839, USA, e-mail: theodorebach@gmail.com

1 Introduction: An Engineering Problem for Social Ontologists

A number of researchers have made clear that the choice between “essentialism” and “eliminativism” with respect to social categories is a false dilemma. Theoretical models and resources are now available – for example Ian Hacking’s “looping effect” and Sally Haslanger’s “discursive constructions” – that provide principled grounds for understanding social categories as both independently real and as the products of contingent social histories. Given the availability of these resources, there is no need to throw out the ontological baby with the essentialist bathwater.

But why flush the essentialism? The reasons at this point are familiar. For one, positing an essence or deep ontological core to a social kind will exclude individuals who *prima facie* belong to that category but who do not possess the essential or core property. Second, to posit essential properties is to theorize current inequities between social groups as natural or essential and therefore justified and unchangeable.¹ Avoiding these essentialist offenses requires rejecting the view that social categories are natural kinds – complexes of properties for which salient properties (e.g. feminine vs. masculine comportment) are ontologically grounded in or causally explained by essential properties (e.g. sexual differences between males and females).

This combination of rejecting natural kind essentialism and affirming the independent reality of social categories presents an engineering problem. Social ontologists require an ontological instrument that can perform simultaneously the jobs of preserving the causal reality of social groups like *men* and *women* while theorizing that reality as grounded in something other than an essential property or natural kind.² Put differently: it is one thing to make clear that a category’s status as “socially constructed” is compatible with its being mind-independently and causally real – it is quite another thing to make clear what degree of ontological unity applies to members of such (essence-lacking) categories.

¹ These problems have several variants that I skip here – but see (Stoljar 1995), (Haslanger 2000), and (Bach 2012).

² The seeming unavailability of such an instrument might explain why at least some researches became invested in the false dilemma described above.

2 “Objective Types” as a Solution to the Engineering Problem

There is growing support for what type of ontological tool is needed to solve the engineering problem: *objective types*. This section explains the ontological category of objective types, contrasts objective types to other ontological categories, and then explores why theorists find objective types to be an attractive solution to the engineering problem.

2.1 What it Means to be Objective, and What it Means to be a Type

Objective types are *types* in the sense that they are sets or classes whose members share a property. Any property will unify a type, and thus any objective property will unify an objective type. Importantly, many objective types are unified on the basis of (objective) relational properties. For example, according to Haslanger (2012), each of the following is an objective type:

- “the things on my desk” (p. 149)
- “things exactly one mile from my dog’s nose” (p. 202)
- “red things” (p. 149)
- “things that are bigger than a bread box” (p. 206)
- “[things] weighing exactly ten pounds” (p. 210)

Objective types are *objective* in the sense that their reality and causal relevance exists independently of our social and cognitive practices. *Biological Centaurs* are not an objective type because they are not real and they have no causal status (they are a type – just not an objective type). On the other hand, “red things” is an objective type because the members of the type – for example individual tomatoes and the surface of Mars – have spatial-temporal and causal properties. Now clearly, social and cognitive practices causally explain the origin and distribution of some red things, for example red fedoras and red lollipops. Nonetheless, the type is objective: if all humans and all conceptual practices were instantly to vanish, red fedoras and red lollipops would continue to exist and causally interact with their surroundings (for example, exerting small amounts of pressure on the hooks of hat racks). So a type’s being “socially constructed” is compatible with it being objective.

2.2 Objective Types Contrasted with Gerrymandered Sets and Natural Kinds

As an ontological category, the members of an objective type exemplify a degree of unity that is above the level of a gerrymandered set but below the level of natural kinds. With respect to the former contrast, Haslanger informs us that there are “highly gerrymandered sets of things that do not have any unity at all and so do not constitute a type” (2012, p. 372, fn. 8). But even here we need to be careful. One might think that a state lottery’s random drawing of lotto balls constitutes a gerrymandered set rather than an objective type. Yet Haslanger acknowledges that this set is an objective type in virtue of its members sharing the objective and relational property *being selected by the lotto machine* (Haslanger 2012, p. 208). That this point generalizes indicates just how minimal the unity requirement is for objective types. For any gerrymandered set, at the moment that one thinks of its members, this very act of cognitive selection will relationally and objectively unify items into an objective type. The implication is that the only possible groupings of mind-independent entities that would *not* count as an objective type would be those potentially gerrymandered groups – groupings that have not, or will never be, actually gerrymandered or grouped.

There is a sharper ontological distinction between objective types and natural kinds. Natural kinds are sets for which members tend to share a syndrome of properties *for the same reason*. There is an ontological ground – call it an essential property – that explains why one member of the kind resembles in various ways other members of the kind. This ontological ground (or essence) is what accounts for the reliable co-projectability of properties. These features of natural kinds explain why they are epistemically and scientifically valuable: you can study one member of a kind and then make grounded, non-accidental generalizations to other members of that kind. For example, biologists can study just a few members of an animal species and then reliably project information to unencountered members of the same species. Chemists can study a bar of gold and then project a syndrome of properties to unencountered deposits. And so on.

Ruth Millikan and Richard Boyd have independently codified highly developed accounts of natural kinds along these lines. Their accounts are similar, and recommending both is that they are situated in more general philosophical systems that explain the interface between human conceptual powers, scientific progress, and mind-independent structures. Applying Millikan’s view (see Millikan 1984, ch. 16; Millikan 1999; Millikan 2000, ch. 2) there are at least three types of ontological ground that support reliable induction. There are “eternal kinds” like gold and water for which the natural laws that govern intrinsic

properties explain likenesses between kind members.³ Then there are “historical kinds,” or natural kinds with historical essences. Examples include *Homo sapiens*, Chevy Novas, McDonald’s, and the lullaby Twinkle Twinkle Little Star. Here, kind members tend to share likenesses due to historically situated copying processes. Third, Millikan claims that individuals – for example Hillary Rodham Clinton, or the Statue of David – are natural kinds. You can study an individual and then make grounded inferences to past and future instances of that individual. The ontological glue grounding induction for individuals includes forms of replication, conservation laws, and homeostatic systems.⁴

In contrast to these three types of natural kinds, objective types lack such a grounded unity. Their members share a property (or properties), but they need not share that property (or properties) for the same reason. In addition, a taxonomy of objective types is one that imposes both empirical and epistemic limits – limits that do not apply for natural kind concepts. This is because objective type concepts *classify* the world, whereas natural kind concepts, when they refer, *identify* kinds and individuals in the world. Concepts that classify rather than identify are limited in the extent to which their users can be wrong about, learn about, and synthesize data about the target. I will return to these important contrasts and their political implications below, in Sections (3.5) and (3.6).

3 Kornblith (1995) provides a succinct discussion of this “eternal” source of property correlation in the case of water: “One may not just slap together protons, neutrons, and electrons in any configuration whatsoever; only certain arrangements will form stable configurations....H₂O is a possible molecule, but HO₂ is not” (Kornblith 1995, p. 36). Had the more or less eternal laws of nature been different – say, if conditions in the Big Bang were slightly different – then perhaps HO₂ rather than H₂O would have been the stable configuration, in which case we would have a distinct natural kind with a distinct syndrome of observable properties.

4 As Millikan explains, “Xavier today is much like Xavier yesterday because Xavier today directly resulted from Xavier yesterday, in accordance with certain kinds of conservation laws, and certain patterns of homeostasis, and because of replications of his somatic cells....there is a deep similarity between individuals and many historical kinds” (Millikan 2000, p. 24). Given this point, there may not always be a strict border between a natural kind’s status as a historical individual and a historical kind. A rule of thumb may be that if conservation forces figure centrally in the explanation of likeness between instances (or time-slices) then one can categorize as a historical individual, and if the kind category grounds generalizations that make it the subject of one of the special sciences then one can categorize it as a historical kind. On this latter point, Griffiths (1999) criticizes the decision of Ghiselin and Hull to categorize species as individuals rather than historical kinds on the basis of their adoption of an incorrect (positivist) model of kinds that requires kinds to support exceptionless generalizations. Finally, even if historical kinds and their associated exception-prone generalizations are often the subject matter of the special sciences, not all historical kinds, for example Twinkle Twinkle Little Star, will be of scientific interest.

2.3 The Theoretical Role of Objective Types in Contemporary Social Ontology

A number of influential philosophers investigating the ontology of social categories agree that paradigmatic and politically important categories refer to objective types. Linda Alcoff offers a theory of gender in terms of “Positionality” – a notion that makes primary an individual’s relationship to reproductive role. Situating this theory within a wider ontological framework, Alcoff explains that “to categorize human beings on the basis of a biological division of reproductive roles is thus to recognize an objective type” (Alcoff 2005, p. 175). Natalie Stoljar contends that “the type ‘woman’ is a type in virtue of the resemblance structure which obtains among individual members of the type. Women constitute a type on the basis of the real (natural and social) similarities among the members of the type...” (Stoljar 1995, p. 282–283). In so far as *resembling a paradigm* is an objective property (and it is, given the discussion from Section 2.2), then Stoljar’s gender types (the flipside of her cluster concepts) are objective types.

In a series of articles Sally Haslanger has argued that men and women are objective types defined by their positions in a social hierarchy. Haslanger provides the most detailed account of “type-objectivism,” as she terms it, and I will examine that position below. Before doing so, it is important to make clear that several other proposals – even if they do not explicitly use the expression “objective type” – occupy and derive shelter within the logical space carved out by this ontological category. For example, Glasgow and Woodward (2015) argue for the objectivity of “groups of people who are distinguished from other groups by having certain visible features (like skin color) to a significantly disproportionate degree” (Glasgow and Woodward 2015, p. 452). To do so, they appeal to what they term “basic racial realism.” They argue that this brand of realism, while admittedly thin, is nonetheless an “intuitively plausible piece of general metaphysics” (Glasgow and Woodward 2015, p. 449). It is supposed to be the same piece of metaphysics that grounds the reality of general categories like “things around the tree.” They are clear that this metaphysical resource lacks natural kind status:

...the metaphysical status of *things around a tree* is not obviously robust. In particular, *things around a tree* is not a natural kind....There are not, for example, many properties that things around trees as such have in common....Nor do they have some ‘theoretically important property’....or explanatory value.... (Glasgow and Woodward 2015, p. 450)

At the same time, they maintain that the metaphysical resource is sufficiently real:

...the things around the tree appear to be real, not just as individuals, but *as* things around the tree – again, things around a tree exist mind-independently as things around a tree. When we say that the things around the tree are real in this way – real in the sense of existing mind independently as such – we maintain that ‘real’ operates as an ordinary-language term. These points suggest that there is a perfectly ordinary sense of ‘real’ in which kinds can be real even when they are not natural or scientifically relevant kinds. (Glasgow and Woodward 2015, p. 450–451)

Substitute the terms “real” and “kind” with “objective” and “type” in the above passages, and it is clear that Glasgow and Woodward’s “basic kinds” are cut from the very same ontological cloth as objective types. Both schemas walk the same line between ontological robustness (or lack thereof) and objective reality.

So why are social theorists so keen on this metaphysical category? As glossed in the introduction, social ontologists are caught in a political dilemma. On the one hand, concerns about justice pressure the social researcher to articulate a realist ontology for social categories. This is because the political problems attending social categories and their occupants are real (as real as “bridge abutments and fists in the face,” in the words of Catharine MacKinnon⁵) and also because effective advocacy requires the affirmation rather than denial of the reality of social groups. On the other hand, there are political dangers if one goes too far in the other ontological direction; the application of robust metaphysical schemas to the social domain have historically generated their own set of political injustices and social inaccuracies that include setting unhealthy normative standards for group membership, modeling certain attributes (e.g. subordination) as natural and/or necessary for group members, and excluding *prima facie* group members.

Objective types, because they are uniquely positioned in logical space in order to balance abstract unity with mind-independent objective status, appear to provide the only viable escape route from this political-ontological dilemma. The reality of social groups, *qua* objective, sufficiently grounds political representation, and yet that reality is not too robust so as to generate further injustices. In short, a theory of objective types is a metaphysical balancing act that doubles as a political balancing act – a virtue that Haslanger (2012) explores in the case of gender and race, and that Glasgow and Woodward (2015) advertise in the case of race.

2.3.1 Objective Type Classifications and Contextual Values

Haslanger has been the most clear about the reasons for applying type-objectivism to the social domain. To understand those reasons, we will need to

⁵ Quoted in (Haslanger 2012, p. 84).

examine Haslanger's more general, value-laden epistemological commitments, as these commitments are central to her argument for the political value of type-objectivism.

These commitments reserve an important, and in many cases necessary, role for contextual values in the sound investigation and classification of the objective world. According to Haslanger, what secures this role for contextual values is an overabundance of worldly facts:

.... an unconstrained search for truth would yield chaos, not theory; truths are too easy to come by, there are too many of them. Given time and inclination, I could tell you many truths – some trivial, some interesting, many boring – about my physical surroundings. But a random collection of facts does not make a theory; they are a disorganized jumble. In the context of theorizing, some truths are more significant than others because they are relevant to answering the question that guides the inquiry. (Haslanger 2012, p. 226)

Because the world bubbles over with structure in this way, what will count as an impartial, complete, and satisfactory theoretical organization of that structure will be relative to the question(s) guiding the theoretical investigation. Anderson (1995) demonstrates this idea in reference to a theory about Jewish involvement in the Atlantic slave trade put forward in the book *The Secret Relationship*. Anderson grants that *The Secret Relationship* contains only empirical truths. Nonetheless, she argues that we should judge the book's theoretical organization of those empirical truths as partial and biased. We should do so by assessing the book's organization of factual claims in reference to the value-laden question that the book purports to investigate, namely, "Do Jews deserve special moral opprobrium or blame for their roles in the Atlantic slave system or bear special moral responsibility for that system's operations?" (Anderson 1995, p. 40). When we do this, claims Anderson, it becomes clear that the book's theory distorts the target phenomena and thereby lacks objectivity. It is only in reference to the contextual values that frame the book's investigation, then, that we are able to assess whether or not significant facts were left out.

Applying this framework to projects in social ontology, Haslanger claims that the legitimacy of a social taxonomy will depend on what questions are driving the taxonomic project in the first place. For Haslanger and others, those questions are rooted in feminist and anti-racist projects: "why might feminist antiracists want or need the concepts of gender and race? What work can they do for us?" (Haslanger 2012, p. 226). Haslanger continues: "At the most general level, the task is to develop accounts of gender and race that will be effective tools in the fight against injustice" (Haslanger 2012, p. 226).

Haslanger offers her hierarchical definitions of gender and race as the tools that we need for that fight.⁶ The hierarchical definitions describe objective types, and purposefully so. Had they described gerrymandered sets or natural kinds, then they would not have been adequate tools in the fight against injustice. Recalling the political dilemma described earlier, gerrymandered sets (or an eliminative approach) lack the objectivity needed for political representation, and natural kinds impose overly strict membership conditions that cause various political problems. By way of the balancing act of type-objectivism, Haslanger's definitions are intended to avoid these political errors. The stipulated classes are objectively real while "...compatible with the kinds of cultural variation that feminist inquiry has revealed, for the substantive content of women's position and the ways of justifying it can vary enormously. *Admittedly, the account accommodates such variation by being very abstract....*" (Haslanger 2012, p. 239; emphasis added). Thus had the taxonomic choice been less abstract – for example, had the taxonomy imposed explanatory essences – then one would pay a political price in terms of (among other things) failing to accommodate politically relevant forms of variation. Glasgow and Woodward (2015) argue similarly, noting that the transition from a "basic" race ontology to a natural kind ontology will bring with it "onerous" commitments.⁷

We can summarize the social appeal of objective types as follows. Begin with the understanding that there is lots of real structure in the world. Given such structure we are justified – compelled, even – in appealing to background goals and values when making taxonomic judgments. Given the specific goals and values of feminism and anti-racism – and evaluating those goals in light of a political dilemma between representation and exclusion – we ought to classify social categories as abstractly defined objective types.

⁶ See Haslanger (2000): "S is a woman iff S is systematically subordinated along some dimension (economic, political, legal, social, etc.), and S is "marked" as a target for this treatment by observed or imagined bodily features presumed to be evidence of a female's biological role in reproduction;" "S is a man iff S is systematically privileged along some dimension (economic, political, legal, social, etc.), and S is "marked" as a target for this treatment by observed or imagined bodily features presumed to be evidence of a male's biological role in reproduction" (Haslanger 2000, p. 50).

⁷ They claim that "basic racial realism secures much of what we want from a good theory of race, without committing to the more onerous ontology of race put forth in constructionism and biological realism" (Glasgow and Woodward 2015, p. 457).

3 Objective Types as a Trojan Horse and the Political Importance of Natural Kinds

In this Section I argue that, from the perspective of the political goals of fighting oppression and marginalization, the category of objective types is a Trojan horse. It looks like a friendly gift, but really it smuggles in several commitments that ultimately frustrate the political goals that motivate using type-objectivism as a taxonomic tool in the social domain. Those goals would be better served, I argue, by modeling social categories as natural kinds.

3.1 Empirical Adequacy as a Condition for Political Adequacy

As discussed in Section (2.3.1), social researchers justify the objective type construal of social categories in reference to the normative questions that guide theoretical investigation. But these same researchers also contend that such investigations must do justice to the empirical phenomena. Haslanger notes that a value-guided theoretical investigation “begins with a rough understanding of the salient facts, and works to construct a conceptual framework that can offer a useful way of organizing them,” and that “in order to assess realistically what tools we need, and why we need them, [value-laden theoretical investigations] depend crucially on descriptive efforts” (Haslanger 2012, p. 353, fn. 22).

Beyond attending to the “salient facts,” just how much freedom is there when selecting theoretical organizations on the basis of contextual values? This question is particularly pressing when the researcher’s normative interests are directed at groups and phenomena that are constrained and shaped by causal forces. For example, researchers who are interested in fighting sexism and racism will want to have theories that organize correctly the social causes of structural oppression, gender and race identity, socialization, and various other causal-explanatory relationships. In such cases, one wonders how the political goal can be met if there is any cleavage at all between the selected theoretical organization and the empirically best theoretical organization.

This tight relationship between the correct tracking of causal-explanatory phenomena and moral outcomes is front and center for at least some social researchers, for example Ian Hacking:

....kinds of people and their behavior which (it is hoped) can enter into laws – laws that if we knew them we would use to change present conditions, and predict what would ensue. We want the right classification – the correct sorting of child abuse or teen-age pregnancy – so that confronted by abusive parents or pregnant teenagers we can embark on a course

of action that will change them for the better and will prevent others from joining their ranks...We want principles according to which we can interfere, intervene, help, and improve. (Hacking 1995, p. 360–361)

Elizabeth Anderson – whose codification of value-laden epistemology, recall, informs several of Haslanger’s commitments – also acknowledges the same general point. She claims that “Empirical adequacy is important not just on epistemic grounds but because an empirically inadequate theory cannot satisfy our practical interests in predicting and controlling phenomena” (Anderson 1995, p. 30).

Our issue now concerns what is required for a theory to be “empirically adequate.” I argue below that if certain conditions are satisfied, then a theoretical classification based on objective types will fail to be empirically adequate. If this is correct, then given that empirical adequacy is a necessary condition for meeting practical goals, then in certain conditions a theoretical taxonomy that is based on objective types will fail to answer to the contextual values (e.g. fighting injustice) that guide the theoretical investigation.

3.2 Conditions for the Empirical Inadequacy, and Thereby Political Inadequacy, of Objective Types

Why should we think that certain theoretical classifications based on objective types might lack empirical adequacy? The idea gets a foothold when we recall that an objective type classification is merely an assertion that a group of items share a property or a set of properties. But in cases where the selected property or properties do not accurately map objectively existing causal-explanatory structures, the social theorist will not have access to forms of understanding and empirical possibilities that bear on politically important features of the target social phenomena. In these cases, the objective type classification will obstruct the fight against injustice because it will limit the social theorist’s ability to predict, control, intervene, and learn about the target phenomena.

To be clear, I am not claiming that objective type classifications are necessarily empirically inadequate. There may be some value-laden investigations that do not depend crucially on tracking particular causal explanatory structures in the world. For example, it is not obvious that the project of determining which taxonomic categories will best serve one’s interests in organizing a vinyl record collection must depend crucially on mapping or tracking causal-explanatory mind-independent structures. This would be the case if one’s goal was to produce a record shelf with a kaleidoscopic array of colors, in which case one could safely overlook real kinds of records, e.g. Stax Records vs Beatles-era Apple Records.

However, investigative projects in the social domain do not generally have this low-stakes quality. As the above quotations from Anderson and Hacking help indicate, value-laden investigations of vulnerable and marginalized social populations *do* depend on the accurate theoretical organization of mind-independent and causal-explanatory structures were these structures, as a matter of fact, to populate one's target domain. So while it may be possible to meet one's record organizing goals while overlooking real kinds of records, it is far less plausible that one could meet the normative goals involved in classifying psychological disorders if one were to overlook real kinds of pathologies.

The discussion above makes reference to real existing kinds. This indicates an additional condition for assessing the empirical inadequacy of objective types. In order for an objective-type taxonomy to be empirically inadequate, it has to be the case that in opting for that classification one is at the same time passing over an alternative classification that would better capture causal-explanatory structure and mind-independent property syndromes in the target domain. Theoretical organizations that might be overlooked in this respect are categories that capture natural kinds. Adding to the discussion from (2.2), we should think of historical and other natural kinds as mind-independent organizations of facts, causal forces, explanatory relationships, and empirical possibilities. While most of logical space is empty (contrary to Locke's speculation about an absence of "gaps" and "chasms"),⁸ there are certain pockets of that space in which causal forces organize and ground property syndromes. In the social domain, these causal mechanisms consist primarily in cultural mechanisms of replication and design (more on this in Section 3.4). These pockets (or "clumps" and "clots," as Millikan 2010 calls them) are what make the world knowable (Kornblith 1995). Were they to obtain in the social domain, accurately identifying them would be politically important.

In flagging the epistemic importance of such mind independent organizations of properties and empirical possibilities, I am not denying that human classifications impact which property syndromes are objectively stable – they clearly do, particularly in the social domain (i.e. Hacking's "looping effect"). But importantly, when classifications do have this type of impact, it is because they are making actual causal contributions to specific cultural mechanisms of replication and design that causally explain likenesses between members of a historical kind. To this extent such patterns of classification are themselves mind-independent empirical processes that, along with the property syndromes that

⁸ See especially Kornblith's discussion of Locke and natural kinds in (Kornblith 1995). Millikan (2010) is also helpful in explaining Locke's error, noting that while there are "well over two million separate extant species of animals...they do not generally fade into each other, let alone into shoe horns and alarm clocks" (Millikan 2010, p. 58).

they help causally maintain, we are in a position to identify or misidentify, be ignorant or knowledgeable about, and so on.

We can formalize these claims in the following schema:

Conditions for the empirical inadequacy, and thereby political inadequacy, of objective type classifications:

A theoretical classification based on objective types is empirically inadequate just in case the following three conditions are all satisfied:

- (a) Accurately tracking causal-explanatory structures is important for achieving the practical goals that guide one's theoretical investigation of the target domain.
- (b) There is theoretical and empirical support for a taxonomy of natural kinds in the target domain.
- (c) Selecting an objective type classification will screen off empirical generalizations and social-political possibilities that would have been otherwise available under the natural-kind taxonomy.

I will assume that the earlier discussion and references to Hacking and Anderson regarding the moral-political importance of accurately identifying and tracking causal regularities and mind-independent property syndromes is sufficient to hold that condition (a) is generally satisfied of the social domain (but I will have a bit more to say about this below). In Section (3.4) I will indicate why there is empirical and theoretical support for a natural kind taxonomy in the social domain and that condition (b) is satisfied. Below, I motivate condition (c) and set the stage for the discussion of natural kinds in Section (3.4).

3.3 Critique of the Argument that Natural Kinds do not Constrain Value-Laden Theory Selection

Anderson (1995) argues that natural kind classifications do not significantly constrain theory selection. Anderson's goal is to limit the degree to which one might appeal to natural kinds in defending specific theoretical classifications. She wants to make sure that there is sufficient room for contextual values during theory justification. In the context of Haslanger's project, this bracketing of the importance of natural kinds clears space for the use of objective type classifications as specified by background feminist and anti-racist values. Thus, any defense of the political importance of using natural kind concepts rather than politically inspired objective type classifications will need to address this critique.

Anderson develops her argument as follows:

...the world is too complex and messy to be organized into a few layers of all-inclusive and mutually exclusive classifications that account for all causal regularities. For each classification that supports some causal regularity, there are likely to be some other crosscutting ones in the neighborhood that bear a causal relation to some other phenomenon. So criteria of epistemic significance alone do not tell us which classifications to base our theory on. (Anderson 1995, p. 45)

To substantiate this idea, Anderson appeals in a footnote to Dupré's defense of "promiscuous realism" (as does Haslanger in several footnotes: Haslanger 2012, p. 188, fn. 8; 2012, p. 91, fn. 21). Anderson also gives the example of *pathogen* as a theoretical classification that, while not tracking a natural kind, responds to contextual health-based interests and is a sound theoretical classification.

Let's start with the pathogen example. The category pathogen groups together any biological item that can cause an infectious disease. I am happy to grant that this category is not a natural kind. In fact, it is an objective type. The unified biological items simply share an objective relational property – causing disease – and they appear to have nothing or little else in common. Unlike the members of a natural kind, there is no univocal reason for why one item in this class resembles in various ways other items in this class. For example, there is no univocal reason for why streptococcus, tapeworm, black bread mold, and PrPSc (a proteinacious infectious particle) each cause disease.

While I agree that the category pathogen is useful for advancing our practical interests, the example does not call into question the conditions for the empirical inadequacy of objective type classifications as described above. This is because in the case of pathogen, condition (c) is not met, and thus the conditions do not give the unintuitive verdict that the classification is empirically and thus politically inadequate. To see why condition (c) is not met, we need to look more closely at the pathogen category. The objective type classification "pathogen" represents a disjunction of natural kinds. Each disjunct in this class – streptococcus, tapeworm, PrPSc, etc. – describes a different natural kind, the members of which (e.g. individual PrPSc particles) *do* share similarities for the same reason. But importantly, in creating this disjunctive category "pathogen," we are not thereby forced to screen off empirical generalizations that are grounded by the natural kind status of the category's disjuncts. For example, by including tapeworm in the category "pathogen," we are not thereby forced to override empirical generalizations about tapeworms that derive from the empirical study of tapeworms *qua* a natural kind (e.g. a historically unified species). In contrast, and as I will describe below, such screening-off does occur in the social domain when social researchers privilege objective type classifications over the identification of natural kinds.

Can type-objectivists find refuge in Dupré's "promiscuous realism" – the claim that there are many legitimate but crosscutting taxonomies of the objective

world? They cannot – at least not in such a way as to rescue them from the proposed conditions of empirical inadequacy. Dupré builds much of his case on the proliferation of species concepts in biological systematics. However, these rival species concepts are not objective type classifications. Rather, they are attempts to model the type of natural kind to which all species kinds belong. As such, rival species concepts represent different construals of which properties are essential to a species and account for the reliable co-projectability of properties across members of a species. According to one species concept that explanatory property is reproductive isolation, according to another it is ecological role, according to another it is genealogical relations, and so on. This debate is thus a type of theoretical competition between natural kind categories. But my proposed conditions for empirical inadequacy apply only to cases in which a natural kind category and an objective type classification are set in competition. They do not apply to “natural kind vs. natural kind” theoretical competitions, or even “objective type vs. objective type” theoretical competitions.

When choosing among competing natural kind categories, Dupré’s brand of pluralism might indeed be the best option. Also relevant here are arguments from Ludwig (2015) about the role of non-epistemic values when making framework-level ontological choices. Nonetheless, I will remain neutral here on whether we should be agnostic, pluralist, selective, or hopeful for future integration in these cases. When choosing among competing objective type classifications – that is, in cases for which there is no empirical and theoretical support for a relevant natural kind category and thus condition (b) is not met – I am inclined to agree that we ought to employ a form of pragmatic pluralism. Haslanger’s discussion of the category “incomplete” in the context of student performance (Haslanger 2012, p. 378) is a likely example. In such cases, there are no extant natural (i.e. historical or eternal) kinds of people, and it is important to bring normative considerations to bear when deciding what kinds of people to bring into existence through our labeling and classificatory practices. But once there are real kinds of people, it is politically important to understand them as the real kinds that they are, even if – especially if – we want to change or eliminate properties of those kinds for various moral and political reasons.

Dupré’s other central argument is based on the claim that folk-classifications often cross-classify scientific classifications. Because folk taxonomic classifications are typically based on objective type categories (Dupré’s frequent discussion of the folk category “lily” is an example), these cases *do* represent the sort of “natural kind-vs-objective type” theoretical standoff that was absent in the species debate. Does this point vindicate the empirical adequacy of objective type classifications privileged by social theorists? It does not. The lily example is similar to the vinyl-record sorting project described earlier. In both cases, it is precisely the low-stakes nature of the investigative project (decorating, hobby

sorting) that permits the investigator to drift away from the domain's natural kinds while still satisfying the interests that frame the investigation. But as discussed earlier, and as also codified in condition (a) for the empirical inadequacy of objective types, when the investigative goal consists in fighting social injustice, that goal cannot be met when investigators dismiss or overlook important causal-explanatory regularities, processes, and/or possibilities in the target domain.

To summarize: I am happy to grant that in cases where several legitimate taxonomic candidates are available (e.g. competing or overlapping natural kind categories) and the stakes are high, there can be reasonable grounds for selecting among those candidates on the basis of contextual values (though as I discuss in Section 3.6, one still should not dictate through empirically mindful but politically inspired definitions the extension and nature of such selected candidates). What the above considerations are meant to show is that in cases where there is support for a natural kind taxonomy and the stakes are high, there are not reasonable grounds for selecting an objective type taxonomy over the legitimate taxonomy on the basis of contextual values. Those values are better served by privileging the natural-kind taxonomy.

3.4 Social Categories as Describing Natural Kinds with Historical Essences

In order for the conditions of empirical inadequacy for objective types to be met, it must be the case that there is empirical and theoretical support for positing natural kinds in the target social domain, and that representing those kinds with theoretical categories will yield probable knowledge and disclose important social-political possibilities. Are there such kinds? Traditionally, researchers have assumed that if there were such kinds they would have to have intrinsic biological essences. For example, with respect to gender, we might construe Baron-Cohen's (2004) account as modeling how biological properties causally ground the reliable co-projectability of gender-related social properties. Baron-Cohen argues that physical differences between males and females – specifically, different types of brain-development that result from differing levels of fetal exposure to androgen – probabilistically grounds a syndrome of social, gender properties that includes empathic versus systematizing thinking styles and a STEM career trajectory versus a care-work career trajectory. This is not the view that I defend here.

While I agree that gender kinds are natural kinds for which there is an essential property that probabilistically explains likenesses among kind members, I do not contend that this ontological ground consists in biological or sexual properties. Instead, the glue that holds together members of a social gender kind is

the same *type* of glue that holds together a species: the replicative and reproductive relations that carve out a lineage of common descent. Contemporary biologists reject the view that the unity of a species consists in members sharing some objective property (type-objectivism) or some essential intrinsic property (natural kinds with intrinsic essences). Contemporary biologists instead model a species' membership conditions on the basis of essential historical relationships that each member bears to the same lineage. They also explain likenesses between species members on the basis of these historical, replicative properties. (If *B* is copied from *A* or made a reproduction of *A*, then *B* and *A* share similarities *because B* is copied or made a reproduction from *A*).

Because copying and reproduction are general processes, non-biological kinds such as artifactual kinds are also historical kinds. So too, I argue in (Bach 2012), are social gender kinds like men and women.⁹ The model developed in (Bach 2012) explains how the replicative processes that take place during an individual's ontogenetic history – for example, the socialization processes that are designed to slot a sexed individual into the historical role of men or women – make that individual a reproduction of ancestral men or women. It is thus through these replicative processes that an individual gains a participatory relation to a historical gender lineage. These replicative processes also explain why gendered individuals tend to share certain behavioral, psychological, and social properties. But note that, by locating the explanatory essence of gender-kinds in wider historical processes, this account allows for individuals who fail to exemplify characteristic gender properties – who fail to be masculine or feminine, or privileged or subordinated – to retain their gender status. The account is thus particularly well suited to handle the concerns about commonality and exclusion that we considered in Section (2.3.1). As we will see below, the account also allows us to understand a future egalitarian gender system as being compatible with the current, historically defined gender kinds *men* and *women*. Finally, and as also elaborated on below, the account has the important epistemic advantage of indicating how our concepts for social categories can be hooked directly onto natural kinds, thereby allowing us to track and learn more about these mind-independent structures. A commitment to objective type classifiers, on the other hand, leaves these politically important epistemic practices either mysterious or impossible.

⁹ See also (Millikan 1999) and (Elder 2004) for general discussions of why social kinds might be natural kinds with historical essences. It is also worth mentioning here that the specific account I provide in (Bach 2012) indicates a general schema that has application to other social kinds, for example race. In particular, one might assimilate features of Alcoff's notion of an "ethnorace" (Alcoff 2000), or McPherson's notion of a "socioancestry" (McPherson 2015), to the non-biological essentialist model of natural kinds with historical essences given in (Bach 2012).

3.5 Type-Objectivism as Screening Off Morally and Politically Relevant Information

With the natural-kind account of gender now on the table, we can better examine how objective type construals of social categories can screen off empirical possibilities for the target social phenomena, thereby impeding the investigator's ability to meet normative goals. These considerations thus further explain how objective type classifications can satisfy the third condition (condition c) for empirical inadequacy.

I will start by articulating a general worry, and then I will offer more specific discussion and examples. When one stipulates that the unity of a target social group consists in members sharing some particular property or set of properties – for example subordination/privileging on the basis of sex – then one is making both an empirical and political gamble. One is betting that the best way to understand the target phenomenon or group, with respect to one's moral and political goals, is on the basis of the stipulated definition of the objective type. The specific gamble is that there is not a better, more explanatory classification of the target group, the tracking and description of which could prove pivotal to satisfying one's political and moral goals. Think of it this way: On the direction of background practical interests, the proponent of type-objectivism imposes or stipulates an organization on the world's facts. But if it turns out that there is a natural kind taxonomy available (mind-independent organizations of facts, causes, explanatory relations, and empirical possibilities), then by imposing an objective-type classification one would be ruling out or overlooking empirical generalizations and possibilities that are likely morally and politically significant.

We can make these concerns more concrete. Consider that Haslanger's apparent reason for advancing her hierarchical definitions of gender categories is that, with oppression and privilege built into the definitions, investigations into gender categories will stay focused on the injustice of women's oppression and the "ones that matter" (Haslanger 2012, p. 240).¹⁰ The concern is that this stipulation has sealed off

10 One wonders if the following, hidden assumption is at work here: that what is most politically central or pressing for a marginalized group should also be theorized as that which is constitutive, or definitive, of that group. This principle should give us pause. Consider a very naïve Civil War doctor who, after one day on the job, has encountered only gangrened feet. The doctor has the investigative goal of rehabilitating soldiers who have gangrened feet. But in order to meet that goal it would be a serious mistake to theorize soldiers' feet as essentially gangrened. This would miss an important fact – important relative to the doctor's normative goals – that follows from the empirical and ontological nature of human feet (e.g. that they often survive the eradication of gangrene). The example helps demonstrate how natural kinds out in the world can constrain the effectiveness of our stipulated taxonomies.

important empirical possibilities and facts – facts that, it might turn out, are important or necessary for meeting the normative goals that prompted the objective type classification in the first place. A particularly important empirical and political possibility that Haslanger’s objective type classification screens off is the mutability of social categories like men and women. If, on the other hand, one were to model men and women (and races, etc.) as natural kinds with historical essences, then such mutability becomes a live empirical possibility (recall that historical lineages can survive the loss of a particular property, for example that of being socially subordinated or socially privileged). Importantly, this descriptive possibility then informs political possibilities; on the historical kind construal (but not the objective type construal), advocating for social justice and gender egalitarianism is compatible with advocating for the preservations of the groups *men* and *women*.

Several points deserve emphasis here. First, the empirical possibility that gender is mutable is not something that the investigator imposes on the world under the direction of his or her background contextual values. Rather, it is a political possibility that emerges through an empirical examination of the social domain’s causal-explanatory categories. Hundreds if not thousands of empirical studies make clear that there is a rich clustering of gender related phenomena for men and for women. Why is it that women smile more than men, that men are more likely to get jaw implants and beard transplants, that women perform better on certain empathizing tasks, and that men have greater access to high-paying positions? The natural kind theorist looks to the world to make sense of all of these regularities. As sketched in Section (3.4), we would do well to explain such property syndromes in a way that is analogous to how we explain property syndromes in both the biological domain and the artifactual domain – that is, in reference to the replicative and reproductive causal processes through which a historical system constructs lineages. But the important point is this: once we are in possession of a natural kind concept that is hooking onto a natural kind, we can use it to learn new things and recognize unanticipated possibilities for the kind that we are tracking. For example, one can use the gender concepts from (3.4) in order to claim on empirical grounds that gender lineages can persist in the absence of injustice. This in turn illuminates a political path for fighting injustice while preserving current gender groups. But such political possibilities come to the surface as a result of tracking the world’s inductively rich categories and not, in the first instance, by tracking our values.

Second, note how this empirical possibility can inform (or re-inform) the value-laden questions that frame the investigative project. Armed with the knowledge that current men and women can persist in an egalitarian gender system, we might ask different questions when framing our investigative project. Instead of asking “which concepts will be effective in the fight against injustice?” we might ask “which concepts will be effective for preserving gender while fighting

injustice?” The stipulative gamble of type-objectivism is that one might seal off questions like the latter question. To privilege an objective kind construal of a social category is thus to place a bet that the property around which one is analytically defining the class is *not* a property that, instead, reliably projects from a more fundamental essential property. Were *that* the case, then different metaphysical and political possibilities present themselves – possibilities that could better inform the normative questions that guide one’s investigation.

Third, there are potentially illustrative examples from other investigative domains of this type of taxonomic gamble. In biological systematics, the taxonomic school of pheneticism grouped together organisms on the basis of overall similarity rather than common descent. This taxonomic method is now discredited and rejected. Nonetheless, we can imagine a value-laden investigation into the protection of a particular endangered species that failed on account of its using a phenetic rather than historical species taxonomy.¹¹ Or consider E.O. Wilson’s reflections on the dangers of misguided biological taxonomies as reported in (Griffiths 1997):

In a discussion of biodiversity E.O. Wilson recounts how attempts to control the mosquitoes that spread malaria were hampered by bad taxonomy of the *Anopheles* genus. A little later he recounts how the modern recognition of many species of the protozoan *Paramecium* creates “a strong temptation to ignore the biological complexity and stay with the three old, easy species, but the malaria example counsels otherwise. Biologists know in their hearts that there can be no compromise on matters of such importance (1992, 46).” (Griffiths 1997, p. 191)

Griffiths’ and Wilson’s point is not that, when the stakes are high, biologists know the correct solution to the species debate. The point is that when the stakes are high, biological taxonomists will (and ought) defer ultimately to what I have been calling “mind independent organizations of empirical facts, processes, and possibilities.” They will not and ought not to defer to a projected class that, as perhaps revealed by further investigation, was an artifact of theorizing.

3.6 Type-Objectivism as Obstructing Semantic and Epistemic Progress

There is an alternative way to argue for the political advantage of modeling social categories as natural kinds rather than objective types. To motivate this alternative,

¹¹ Along these lines, see the discussion in (Daugherty et al. 1990) about how the failure to recognize a morphologically similar but genealogically distinct subspecies of the Tuatara lizard led to its unnecessary extinction.

we can first consider a possible response from the type-objectivist to the foregoing critique. The type-objectivist might claim that they are free to pack into their class definition additional axes of similarity and that this will achieve greater explanatory power. For example, instead of focusing on just the abstract structural similarity of subordination/privileging, or just the superficial visible similarities that unify basic race kinds (Glasgow and Woodward), they can appeal to many additional shared properties that are of political interest. Moreover, they could remind us that natural-kind categories are themselves of varying explanatory pay-offs.¹²

In reply, we should first inquire about the justification for the inclusion of the specific properties that now enrich the objective type classification. When such a story is fleshed out, it will likely be the case that the type's explanatory power derives from its approximation of a natural kind category that is doing the actual explanatory work, in which case we ought to replace the parasitic classification with the more accurate natural kind category.¹³ Of course, replacing the type category with a kind category will require following the current recommendation of deferring to mind-independent natural kinds rather than our theoretical classifications.

Second, and related, the issue of accuracy and not the degree of explanatory robustness is what is most central to the political tradeoff between natural kind essentialism and type-objectivism. By "accuracy" I mean correctly identifying and gathering information over time about the (explanatory) kinds that in fact populate the target domain (assuming condition b.). Such accurate identifications and trackings are important even if the domain's kinds are, by their own nature, "fuzzy" or "rough" such that they afford only few or somewhat probable generalizations. In order to understand those kinds we will need to defer epistemically to their actual borders, quality, and ontological ground – we risk inaccuracies if we instead defer to our own classificatory attempts.

These considerations bring into focus a politically important semantic contrast between natural kind essentialism and type-objectivism. Objective type classifications are empirically informed armchair reflections about politically important groupings. As such, their extensions are simply functions of their

¹² For example, Millikan (2000, p. 24–26) discusses "kinds of betterness and worseness in substances," and Griffiths (1999, p. 216) allows that mere above chance predictions are evidence that a category is tracking a natural kind.

¹³ Griffiths captures the concern well: "there are some categories without causal homeostatic mechanisms, and these should be rejected. These need not be categories which are entirely useless for explanation and induction. Categories that need to be eliminated are more usually ones whose apparent causal homeostasis results from their approximating some other category with real causal homeostasis. Replacing the crude approximation with the category that is doing the actual work is a pure epistemic gain" (Griffiths 1999, p. 216).

defining property or, in the case of multiple properties, the union of those properties; they are analytically but not synthetically projectable.¹⁴ But this concession to the armchair and the implicit commitment to classifiers cuts off access to politically advantageous semantic resources, specifically (and assuming condition b.) a pure, full-throated form of natural kind externalism.¹⁵ Rather than classify, concepts for natural kinds, when they are successful, latch onto structures in the world (often historical kinds and individuals) about which one can learn many new things. Such natural kinds serve as anchors for conceptual change. While it is typical to at first latch onto such kinds in an epistemically limited way, for example with incomplete descriptions and limited methods for re-identification, over time one learns to better describe and re-identify the kind.

With respect to the social domain, this semantic contrast indicates a political cost for type-objectivism in terms of that view's potential to obstruct our ability to learn about social kinds. On the natural kind approach discussed in (3.4), terms like "men" and "women" are identifiers, and we are constantly updating and revising our descriptions and theories of the identified kinds. On the objective type approach to social categories, on the other hand, terms like "men" and "women" function as classifiers, and the act of stipulating the unity for that class – however robust – will set strict limits on what new information one can gain about that class.¹⁶

14 In being told that something is a "red thing," nothing follows other than that it is red and a thing. If the classification expands to include twelve additional properties, nothing follows other than that it is a red thing with these twelve additional properties. As Millikan puts it, such classifications "wear their analytical nature on their sleeve" (Millikan 2000, p. 39). See Millikan (2000) for general discussion of the contrast between classifiers and identifiers.

15 As Fumerton (1989) and Millikan (2000) make clear, several standard versions of "externalism" (e.g. of the variety commonly attributed to Kripke and Putnam) are still committed to a form of descriptivism or "conceptualism," which is the idea that our thoughts and theories about what we are thinking about constrain what we are in fact thinking about. Any such deferral to our theories rather than actual kinds, should they exist, commits the type of semantic political gamble discussed here. On the other hand, Millikan's (2000) account of how concepts get their extension is free of any such descriptivist commitments.

16 Haslanger occasionally appeals to what she calls "objective-type externalism," which is her application of semantic externalism to objective types. She might intend that application to handle at least some of the concerns raised here about conceptual change. Partly on the basis of what I have been claiming about the epistemic contrast between natural kinds and objective types, as well as my claims about the distinction between identifiers and classifiers, I am not convinced by Haslanger's brief argument for the generalization of semantic externalism (particularly the full-throated version discussed here) to objective types (especially objective type classifications that result from normative inputs). Nor do I address here Haslanger's very different and programmatic remarks about reference and meaning raised in her essay "Language, Politics, and 'The Folk': Looking for 'The Meaning' of 'Race'". Perhaps that is a discussion for a different essay.

A related political cost concerns the obstruction of our ability to communicate and exchange knowledge about politically important social kinds. Natural kinds often do not present all or many of their properties to us on a given encounter; different properties might manifest under different conditions. For political reasons, we want the semantic and epistemic flexibility to say that we are talking about the *same* kind even if we are using very different modes (both within the same individual and between individuals) of re-identifying that same kind. This point extends to social scientific research programs that focus on different property ranges for the same kind and thus use different ways of re-identifying and gathering information about that kind. For example, one research program might re-identify the kind *men* through investigations into the type and amount of domestic labor that men engage in, and another research program might re-identify the kind *men* through investigations into men's increasing interest in certain forms of elective cosmetic surgery (e.g. calf implants and beard transplants). The point generalizes to other social kinds (consider women who may not be subordinated, or members of a race who do not sufficiently exemplify Glasgow and Woodward's visible race features). In this epistemic situation it would be politically costly to adopt a taxonomic approach that precludes the semantic resources needed for a fruitful exchange of information between investigations of the same kind. A distinct political cost would result if in order to make room for the possibility that these investigations are about the same thing, one was forced to misdescribe the identified kind as biological (sex) rather than social (gender), in which case one gets wrong the actual source (the ontological ground, the kind essence) of the property correlations.

In conclusion, it is difficult to reconcile the objective type construal of social categories with the possibility that we can learn new and surprising things about the category, or that we might revise or even abandon our conception of the category, or that we can fruitfully synthesize information about disparate instances of the category. At the very least there is a tension here. When type-objectivism stipulates for political reasons who it wants to talk about, it is dictating the extension of that social category. But what if the world suggests otherwise? What do we say in cases where we have compelling evidence that a man is not privileged and that a woman is not subordinated? To stick to one's stipulative guns – in accordance with the analytic projectability of objective types – one could say that such cases indicate distinct objective types. Or – and this is the option that I recommend – one could say that one's objective type classification was approximating a natural kind category all along – that one is learning how additional properties cluster together alongside the original property – that there are good reasons for why we find these properties clustered together in nature – and that we as investigators are open to revising what we know about these kinds as we continue to track them in the social domain. Such a view is consistent with – indeed supportive

and instrumental to – the moral imperative of focusing on and vigorously fighting various forms of social injustice.

4 Conclusion

It is common for researchers to argue on normative grounds that social categories are objective types. This essay argued that a closer examination of the ontological category of objective types, in the context of a set of conditions for empirical adequacy, reveals that social researchers ought to prefer theoretical categories based instead on natural kinds with historical essences. I argued that they ought to form this preference on the basis of the very normative reasons that framed their preference for objective types.

Acknowledgments: I developed this essay in the context of a graduate seminar on social ontology that I taught at Bowling Green State University in the fall of 2015. Many thanks to the excellent graduate students – Natalie Cote, Lisa Gore, John Luchon, Austin McGrath, Elizabeth McPherson, and Amitabha Palmer – for helpful comments and sharp discussion. I am also grateful to two anonymous referees from the Journal of Social Ontology for excellent comments.

Bibliography

- Alcoff, Linda (2000): “Is Latina/o Identity a Racial Identity?”. In: Jorge J. E. Gracia and Pablo De Greiff (Eds.): *Hispanics/Latinos in the United States: Ethnicity, Race, and Rights*. New York: Routledge, p. 23–44.
- Alcoff, Linda (2005): *Visible Identities: Race, Gender, and the Self*. New York: Oxford University Press.
- Anderson, Elizabeth (1995): “Knowledge, Human Interests, and Objectivity in Feminist Epistemology”. In: *Philosophical Topics* 23. No. 2, p. 27–58.
- Bach, Theodore (2012): “Gender Is a Natural Kind with a Historical Essence”. In: *Ethics* 122. No. 2, p. 231–272.
- Baron-Cohen, Simon (2004): *Essential Difference: Male and Female Brains and the Truth about Autism*. New York: Basic Books.
- Daugherty, C. H., A. Cree, J. M. Hay, and M. B. Thompson (1990): “Neglected Taxonomy and Continuing Extinctions of Tuatara (*Sphenodon*)”. In: *Nature* 347, p. 177–179.
- Elder, Crawford (2004): *Real Natures and Familiar Objects*. Cambridge, Mass: Massachusetts Institute of Technology Press.
- Fumerton, Richard (1989): “Russelling Causal Theories of Reference”. In: C. Wade Savage and C. Anthony Anderson (Eds.): *Rereading Russell: Essays in Bertrand Russell’s Metaphysics and Epistemology* 12. University of Minnesota Press, p. 108–118.

- Glasgow, Joshua and Jonathan M. Woodward (2015): "Basic Racial Realism". In: *Journal of the American Philosophical Association* 1. No. 3, p. 449–466.
- Griffiths, Paul E. (1997): *What Emotions Really Are: The Problem of Psychological Categories*. Chicago: University of Chicago Press.
- Griffiths, Paul E. (1999): "Squaring the Circle: Natural Kinds with Historical Essences". In: Robert A. Wilson (Ed.): *Species: New Interdisciplinary Essays*. Cambridge, Mass.: Massachusetts Institute of Technology Press, p. 209–228.
- Hacking, Ian (1995): "The Looping Effects of Human Kinds". In: Dan Sperber, David Premack, and Ann James Premack (Eds.): *Causal Cognition: A Multidisciplinary Debate*. Oxford: Clarendon Press, p. 351–394.
- Haslanger, Sally (2000): "Gender and Race: (What) are They? (What) do we Want Them to Be?". In: *Noûs* 34. No. 1, p. 31–55.
- Haslanger, Sally (2012): *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press.
- Kornblith, Hilary (1995): *Inductive Inference and its Natural Ground: An Essay in Naturalistic Epistemology*. Cambridge, Mass.: Massachusetts Institute of Technology Press.
- Ludwig, David (2015): "Ontological Choices and the Value-Free Ideal". In: *Erkenntnis*, 1–20. doi 10.1007/s10670-015-9793-3.
- McPherson, Lionel K. (2015): "Deflating 'Race'". In: *Journal of the American Philosophical Association* 1. No. 4, p. 674–693.
- Millikan, Ruth (1984): *Language, Truth and Other Biological Categories*. Cambridge, Mass.: Massachusetts Institute of Technology Press.
- Millikan, Ruth (1999): "Historical Kinds and the 'Special Sciences'". In: *Philosophical Studies* 95. No. 1, p. 45–65.
- Millikan, Ruth (2000): *On Clear and Confused Ideas: An Essay about Substance Concepts*. Cambridge: Cambridge University Press.
- Millikan, Ruth (2010): "On Knowing the Meaning; With a Coda on Swampman". In: *Mind* 119. No. 473, p. 43–81.
- Stoljar, Natalie (1995): "Essence, Identity, and the Concept of Woman". In: *Philosophical Topics* 23. No. 2, p. 261–293.