Race as a Regulatory Kind: Constructivist Realism Beyond Naturalism

Mark Brewer

April 2025

Abstract

Despite widespread scientific agreement that human biological diversity is real, the question of whether "race" corresponds to a natural kind remains deeply contested. While some philosophers and scientists continue to explore ways of biologically grounding racial categories, this paper argues that the project of racial naturalism—whether in its essentialist or reformulated variants—remains conceptually, empirically, and metaphysically untenable. Yet this is not a rejection of the reality of race. Rather, I contend that race is a real and powerful social construct, historically forged and materially entrenched, but not a natural kind in the biological or taxonomic sense.

The paper's central aim is twofold. First, it critically dissects both Standard Racial Naturalism (SRN) and New Racial Naturalism (NRN), demonstrating that neither satisfies the conditions required of a scientifically coherent racial taxonomy. Second, it advances a constructive alternative: a biologically informed form of constructivist realism that understands race as a socially constructed, causally efficacious, and materially instantiated kind. To support this framework, I introduce and develop the concept of *regulatory kinds*: social kinds stabilized not merely by looping classification, but by formal institutions, epistemic infrastructures, and administrative mechanisms that actively enforce kindhood. Unlike natural kinds, regulatory kinds derive their coherence from recursive regulation, classification pressure, and material uptake. Race, I argue, is best understood as a regulatory kind—ontologically real, epistemically powerful, and normatively charged.

By framing race as a regulatory kind, I clarify the collapse of NRN and provide a more precise ontological model of race's durability in classificatory regimes. This account preserves the insights of constructivist realism while extending its explanatory power in epistemically and institutionally entangled contexts.

Keywords: race, natural kinds, constructivism, taxonomy, clustering, regulatory kinds, philosophy of biology

Introduction

There is now widespread agreement that racial categories, as traditionally conceived, do not track biological kinds. The notion that races are defined by shared, stable, and biologically

essential properties has been decisively rejected on both empirical and philosophical grounds. However, some philosophers and scientists continue to hold that race may yet be vindicated as a natural kind of some revised form—typically one indexed to population structure, probabilistic clustering, or forensic inference. These so-called "New Racial Naturalist" (NRN) accounts differ from Standard Racial Naturalism (SRN) in rejecting typological assumptions while preserving the hope that race, properly understood, reflects biologically informative divisions.

This paper challenges the plausibility of NRN on both methodological and metaphysical grounds. While its proponents often invoke clustering algorithms, reference populations, or forensic DNA techniques to support the reality of racial kinds, I argue that these classificatory outputs are shaped by inputs already structured by policy-defined social categories. The result is a form of epistemic circularity. Rather than revealing biologically grounded kinds, these practices reproduce pre-existing social classifications under the guise of empirical refinement.[1]

The collapse of NRN, however, does not entail eliminativism about race. Constructivist realism remains a viable framework—one that understands race as a socially constructed, institutionally entrenched, and materially consequential kind. Yet even within this framework, an important question remains underdeveloped: how do racial categories come to exhibit the stability, coherence, and classificatory resilience typically associated with natural kinds?

To answer this, I introduce the concept of a *regulatory kind*: a socially constructed kind whose apparent empirical stability is maintained through recursive classification practices embedded in formal institutions. Regulatory kinds do not track natural properties, but they can mimic the epistemic profile of natural kinds by virtue of their role within administrative, scientific, and policy infrastructures. The regulatory kind schema, developed in the final section of the paper, explains how social inputs (S) are processed through institutional structures (I) and classification procedures (P), generating outputs (O) that reinforce the stability of the kind through recursive uptake.

The paper proceeds as follows. Section 1 distinguishes SRN from NRN, and critically assesses the conceptual and methodological commitments of the latter. Section 2 outlines the appeal and limitations of constructivist realism as a non-eliminativist alternative. Section 3 develops the regulatory kind schema and shows how it clarifies both the failure of NRN and the ontological structure underpinning the empirical persistence of race.

[1] This concern is particularly acute in cases where reference populations in forensic and biomedical research are drawn from census or passport categories, and then fed back into algorithms which appear to re-identify them as stable clusters. For example, STRUCTURE's often-cited $k = 5$ output, which appears to align with the U.S. Office of Management and Budget's racial categories, is frequently treated as empirical support for NRN accounts, despite the circular dependencies involved.

1 Standard Racial Naturalism and the Problem of Essentialism

1.1 Historical and Conceptual Foundations of SRN

Standard Racial Naturalism (SRN) has its intellectual roots in early modern biological taxonomy, where thinkers such as Carl Linnaeus and Johann Friedrich Blumenbach first attempted to classify human populations into distinct races. These efforts were shaped by both the taxonomic ambitions of 18th-century natural science and the sociopolitical imaginaries of European colonial expansion. Linnaeus's 1758 *Systema Naturae* listed four racial groupings based on continent and temperament, while Blumenbach's 1775 dissertation introduced a five-race scheme grounded in craniometry. Though conceptually unsound by modern standards, these classificatory systems established a durable template: race as a natural division within the human species.

Throughout the 19th and early 20th centuries, the concept of race became intertwined with essentialist assumptions. Human variation was interpreted not as a continuous or overlapping distribution of traits, but as a set of discrete, heritable differences marking stable subspecies. The racial categories employed by anthropologists and biologists during this period were presented as taxonomically real, and often correlated with hierarchies of intelligence, temperament, and moral character. What unified these various formulations was the idea that race reflected underlying biological essences—immutable properties that explained both physical and behavioral attributes.

By the mid-20th century, this essentialist framework began to face growing challenges. Advances in genetics, evolutionary theory, and physical anthropology revealed the complexity and fluidity of human variation. Nevertheless, SRN persisted in a reformulated guise. Instead of asserting full-blown essences, post-war defenders of SRN often adopted a populationist stance: races were conceived as statistically definable groups marked by clusters of traits, rather than fixed intrinsic properties. This population-level reframing sought to preserve the scientific legitimacy of race while avoiding the most overt metaphysical commitments of essentialism.

A key moment in this trajectory was Ernst Mayr's articulation of the biological species concept. While Mayr himself used terms like "geographic race" and "subspecies" interchangeably, his broader work signaled a transition away from typological thinking and toward an evolutionary understanding of species. However, the conceptual ambiguity between population structure and taxonomic rank left open a space for SRN to persist. It remained possible, under Mayr's influence, to describe human groups as "geographic races" while retaining the assumption that these groups had some natural kind status.

This historical entanglement between taxonomy, essentialism, and the politics of human difference is what gives SRN its rhetorical and intuitive appeal. Its categories have long been stabilized through institutional practices—census forms, immigration law, forensic databases—such that race appears not merely as a social identity, but as a natural one. SRN draws its durability from this sedimented history: a hybrid of Enlightenment science, typological metaphysics, and bureaucratic normalization. To critically evaluate SRN, then, we must not only test its biological assumptions, but also dislodge its conceptual inheritance.

1.2 The Rejection of Biological Essentialism

A central vulnerability of Standard Racial Naturalism lies in its metaphysical dependence on essentialism. In its classical form, biological essentialism maintains that species—and by extension, subspecies or races—are defined by fixed, internal properties shared by all members. These properties are assumed to determine not only morphology, but also behavior, intelligence, and temperament. Within SRN, essentialism grounds the intuition that race is a deep biological category: that members of a racial group are alike in virtue of some stable, biologically inherited essence.

This view, while once dominant in pre-Darwinian taxonomy, has been largely abandoned in contemporary evolutionary biology. As Okasha (2002) and Sober (1994) have shown, essentialist species concepts have been replaced by population-based models that emphasize dynamic, historical, and relational properties. Species are now typically understood as reproductively isolated populations, lineages with common descent, or ecological clusters—not as groups defined by shared essences. No single trait, or set of traits, is both necessary and sufficient to define a biological species.

This shift has profound implications for SRN. If biological taxonomy no longer supports essentialist thinking at the species level, it is highly implausible that essentialism could underwrite classifications at the subspecies or race level. Indeed, the very notion of "subspecies" in humans becomes suspect if it cannot be grounded in a robust metaphysical framework. The modern understanding of species undermines the idea that race reflects biologically distinct and coherent subdivisions within Homo sapiens.

Moreover, the rejection of essentialism is not merely conceptual—it is also empirical. Human phenotypic and genotypic variation is continuous and overlapping. Traits such as skin color, facial structure, or genetic markers exhibit clinal variation across geographic space, rather than forming discrete clusters. There are no biologically sharp boundaries between populations; any attempt to draw them reflects external interests rather than internal necessity.

To reinforce this, it is worth recalling that most traits associated with race—such as skin pigmentation or hair texture—are influenced by a relatively small number of genes, many of which exhibit significant variation even within local populations. No trait, or set of traits, is universally present in all members of a so-called racial group. The idea of race as a biologically homogeneous kind collapses under scrutiny.

Furthermore, evolutionary processes such as gene flow, genetic drift, and local adaptation generate complex patterns of human variation that defy simple classification. High levels of admixture, both historical and contemporary, ensure that no population is genetically isolated or static. This is not merely a technical point: it undermines the very logic by which SRN seeks to naturalize race.

It is also important to emphasize the epistemic consequences of rejecting essentialism. If racial categories do not reflect real natural divisions, then their continued use in scientific discourse requires strong justification. To invoke race as a biological kind is to make a metaphysical commitment that must be supported by more than statistical convenience or historical inertia. Without essentialism, SRN lacks the conceptual machinery to justify that commitment.

In sum, the rejection of biological essentialism removes the conceptual foundation on

which SRN depends. Modern biology recognizes that species are dynamic, overlapping, and historically contingent—characteristics that render the notion of fixed racial essences untenable. If essentialism cannot support taxonomic classification at the species level, then it cannot rescue SRN from conceptual collapse.

1.3 The First Hurdle: From Clusters to Kinds

A central challenge for defenders of Standard Racial Naturalism lies in the slippage between statistical clustering and natural kind classification. Even when patterns of genetic similarity appear across populations, it does not follow that these patterns carve nature at its joints. I call this conceptual gap the *First Hurdle*: the point at which racial naturalists must demonstrate that groupings based on biological data are not merely statistically useful, but taxonomically salient.

This distinction is too often obscured in both scientific and philosophical discussions of race. Genetic clustering, made possible by techniques such as principal component analysis and STRUCTURE algorithms, reveals broad population-level patterns correlated with geography and ancestry. Yet these patterns reflect probabilistic tendencies, not strict boundaries. Clusters emerge from algorithmic settings—how many groups the software is instructed to find, how it weights allele frequencies, and what reference populations are used. They are artifacts of modelling choices, not necessarily of deep biological divisions.

The classic empirical reference point is Lewontin's (1972) finding that roughly 85% of human genetic variation occurs within populations traditionally labelled as races, with only a small percentage accounting for between-group differences. Edwards (2003) responded that while this is true, it does not preclude the existence of correlations across loci that allow for accurate classification. Indeed, clustering techniques can predict self-identified race in U.S. populations with a degree of reliability.

But this statistical success does not settle the metaphysical question. As Hochman (2013) rightly argues, the existence of clusters does not entail the existence of kinds. Clusters tell us where similarities fall, not whether those similarities are essential, causally explanatory, or taxonomically decisive. To move from clusters to kinds is to make a metaphysical leap that demands philosophical and biological justification—not just statistical adequacy.

To illustrate the problem, consider the construction of a hypothetical category based on eye colour—*oculos*. Suppose this trait shows geographic clustering, is heritable, and can be predicted with high accuracy using genomic data. Even if such a group exhibits tight statistical coherence, it would not follow that *oculos* constitutes a biologically meaningful kind. The category lacks explanatory depth: it does not track an underlying biological mechanism or evolutionary process. Its salience is superficial.

The same logic applies to racial clustering. While certain phenotypic and genotypic traits correlate with geographic ancestry, these traits do not necessarily signal shared evolutionary pathways, ecological function, or developmental architecture. The traits selected for clustering—skin pigmentation, hair form, facial morphology—are chosen because they are historically visible and socially salient, not because they reflect deep taxonomic structure. Their coherence is retrospective and contingent, not indicative of natural divisions.

Moreover, genetic admixture undermines any attempt to draw stable boundaries between racial groups. Most human populations exhibit gene flow across generations and geographic

regions, resulting in fuzzy, overlapping genetic signatures. This is not a problem for understanding human diversity—it is a problem only if one seeks to impose rigid taxonomic schemes onto it. Clustering becomes a descriptive tool, not a basis for ontological commitment.

In this context, the First Hurdle remains unmet. Racial groupings produced through statistical methods may be empirically robust for certain purposes, such as ancestry inference or epidemiological risk estimation. But such utility does not confer natural kind status. To assert that race is biologically real in a naturalistic sense requires more than pattern recognition—it requires explanatory depth, metaphysical justification, and biological coherence. SRN provides none of these.

1.4 Thought Experiments and Illustrative Cases

To further test the conceptual integrity of SRN, it is useful to examine a range of analogical and hypothetical cases that mirror its structure but lack its rhetorical weight. These cases allow us to isolate the principles at work—particularly the temptation to infer kindhood from clustering, and the tendency to confuse visibility with explanatory significance.

We have already seen the basic structure of the *oculos* thought experiment: a hypothetical classification based on eye colour that, while biologically real and heritable, lacks deeper taxonomic meaning. The same reasoning can be applied to a range of other phenotypic traits that exhibit population-level trends: ear shape, hair texture, blood type, or metabolic tolerance to specific enzymes.

Consider lactose tolerance. The ability to digest lactose into adulthood is influenced by a genetic variant (most commonly the LCT gene), which appears with higher frequency in populations with a history of pastoralism. This trait shows strong geographic clustering and clear heritability. Yet we do not posit the existence of a "lactose-tolerant race" or regard lactose metabolism as taxonomically significant. Why? Because the trait lacks broader explanatory scope. It is one biological feature among many, not a marker of biological kindhood.

Similarly, resistance to malaria through sickle-cell traits (HbS allele) occurs in several geographically dispersed populations in Africa, the Middle East, and parts of India. This trait is strongly selected for in malarial environments, and its presence can be reliably predicted based on environmental ancestry. Yet again, we do not interpret the presence of the HbS allele as indicative of a racial boundary. It is a local adaptation to a specific ecological pressure, not a taxonomic division.

These analogies demonstrate that heritability, frequency, and even functional significance do not justify the inference to natural kinds. What matters is whether the trait or trait complex participates in a larger pattern of evolutionary differentiation that would warrant classification. Most traits associated with racial classification—such as skin pigmentation or craniofacial morphology—fail to meet this standard. They are polygenic, subject to environmental modulation, and widespread across populations.

Moreover, these examples show that clustering and coherence do not necessarily entail conceptual unity. Many biological categories exhibit statistical patterning without forming robust kinds. The temptation to treat any such pattern as evidence of kindhood is a

residue of typological thinking—an intellectual reflex that modern biology has worked hard to overcome.

The racial categories posited by SRN survive in large part because of their social visibility. Traits like skin colour and facial structure have historically been used as visual markers of group membership, reinforced by cultural, political, and institutional practices. These features make racial categories appear natural even when their biological grounding is weak. By contrast, categories based on metabolic enzymes or blood antigens, though biologically precise, lack intuitive salience and thus do not acquire the same kind of metaphysical inertia.

The thought experiments and analogies presented here are not meant to trivialize human variation. Rather, they are designed to show that the mere existence of statistically coherent traits, even when heritable and adaptive, does not justify the assertion of taxonomic kinds. The burden remains on SRN to show that racial categories do more than summarize superficial features—that they track explanatory divisions within biological systems. To date, that burden has not been met.

1.5 From Pattern to Kind: The Illusion of Taxonomic Depth

To further test the conceptual integrity of SRN, it is useful to examine a range of analogical and hypothetical cases that mirror its structure but lack its rhetorical weight. These cases allow us to isolate the principles at work—particularly the temptation to infer kindhood from clustering, and the tendency to confuse visibility with explanatory significance.

We have already seen the basic structure of the *oculos* thought experiment: a hypothetical classification based on eye colour that, while biologically real and heritable, lacks deeper taxonomic meaning. The same reasoning can be applied to a range of other phenotypic traits that exhibit population-level trends: ear shape, hair texture, blood type, or metabolic tolerance to specific enzymes.

Consider lactose tolerance. The ability to digest lactose into adulthood is influenced by a genetic variant (most commonly the *LCT* gene), which appears with higher frequency in populations with a history of pastoralism. This trait shows strong geographic clustering and clear heritability. Yet we do not posit the existence of a "lactose-tolerant race" or regard lactose metabolism as taxonomically significant. Why? Because the trait lacks broader explanatory scope. It is one biological feature among many, not a marker of biological kindhood.

Similarly, resistance to malaria through sickle-cell traits (HbS allele) occurs in several geographically dispersed populations in Africa, the Middle East, and parts of India. This trait is strongly selected for in malarial environments, and its presence can be reliably predicted based on environmental ancestry. Yet again, we do not interpret the presence of the HbS allele as indicative of a racial boundary. It is a local adaptation to a specific ecological pressure, not a taxonomic division.

These analogies demonstrate that heritability, frequency, and even functional significance do not justify the inference to natural kinds. What matters is whether the trait or trait complex participates in a larger pattern of evolutionary differentiation that would warrant classification. Most traits associated with racial classification—such as skin pigmentation or craniofacial morphology—fail to meet this standard. They are polygenic, subject to environmental modulation, and widespread across populations.

Moreover, these examples show that clustering and coherence do not necessarily entail conceptual unity. Many biological categories exhibit statistical patterning without forming robust kinds. The temptation to treat any such pattern as evidence of kindhood is a residue of typological thinking—an intellectual reflex that modern biology has worked hard to overcome.

The racial categories posited by SRN survive in large part because of their social visibility. Traits like skin colour and facial structure have historically been used as visual markers of group membership, reinforced by cultural, political, and institutional practices. These features make racial categories appear natural even when their biological grounding is weak. By contrast, categories based on metabolic enzymes or blood antigens, though biologically precise, lack intuitive salience and thus do not acquire the same kind of metaphysical inertia.

The thought experiments and analogies presented here are not meant to trivialize human variation. Rather, they are designed to show that the mere existence of statistically coherent traits, even when heritable and adaptive, does not justify the assertion of taxonomic kinds. The burden remains on SRN to show that racial categories do more than summarize superficial features—that they track explanatory divisions within biological systems. To date, that burden has not been met.

1.6 The Misuse of Forensic and Genomic Tools

In recent years, Standard Racial Naturalism has found renewed rhetorical support through the emergence of forensic and genomic technologies. Tools such as forensic DNA phenotyping (FDP), biogeographic ancestry testing, and admixture analysis appear to lend scientific legitimacy to racial classification by associating genetic data with visible traits or geographic origins. Proponents of SRN often cite these tools as evidence that race has a stable biological foundation, capable of being measured, predicted, and mapped. Yet this interpretation reflects a deep misunderstanding of both the capabilities of these tools and the epistemic status of the categories they appear to reinforce.

Forensic DNA phenotyping is perhaps the most visible of these applications. FDP seeks to reconstruct phenotypic characteristics—such as skin pigmentation, eye colour, or hair texture—from genetic data. These characteristics are then often used to make probabilistic inferences about an individual's ancestral background or racial identity. While the technology can yield predictive accuracy for certain traits, its interpretation is inherently entangled with prior social categories. The phenotypes it predicts are precisely those that have historically been racialized. As Sankar (2012) observes, FDP does not discover race so much as reproduce it—repackaging socially meaningful traits in scientific language.

The circularity here is epistemically fatal. The categories used to train FDP models are themselves based on socially constructed racial labels, often drawn from census data or self-identification. These categories are then used to validate predictions, creating the illusion that race has been biologically located when, in fact, it has merely been operationalized. What results is a form of technological reification: a system that encodes social assumptions into genetic models and then reads them back as biological facts.

Similar concerns apply to ancestry testing and admixture analysis. Companies and researchers often report genetic ancestry in terms that map loosely onto continental racial categories—"Sub-Saharan African," "East Asian," "European," and so on. These labels

are statistically derived from reference populations and bear little resemblance to taxonomic classifications in evolutionary biology. The clusters they produce are shaped by both methodological choices and historical contingencies—who is sampled, how clusters are defined, and what baseline populations are assumed.

While these techniques may serve practical or commercial purposes, they do not support the metaphysical claims of SRN. At best, they offer probabilistic insights into patterns of shared genetic markers across populations. At worst, they conflate genetic similarity with racial kindhood, reinforcing categories that are neither biologically discrete nor taxonomically meaningful.

Moreover, such tools often operate in the absence of clear philosophical criteria for natural kinds. They rely on intuitive judgments about similarity, visibility, and coherence, rather than principled commitments to explanatory depth or biological salience. In doing so, they blur the boundary between descriptive utility and ontological commitment. They may tell us where people come from, but they do not tell us what kinds of beings they are.

Finally, the political stakes of these technologies should not be overlooked. By lending scientific legitimacy to racial categories, forensic and genomic tools risk entrenching existing social hierarchies under the guise of objectivity. They extend the reach of SRN by embedding its assumptions into databases, algorithms, and institutional practices. In this way, the failure of SRN is not just epistemic, but ethical: it misleads inquiry, misrepresents human diversity, and legitimizes social divisions through flawed biological reasoning.

For these reasons, the forensic and genomic revival of racial classification cannot rescue SRN. These tools do not uncover a hidden biological reality; they reformat social taxonomies in genomic terms. They pass the appearance of scientific rigour, but not the test of philosophical or biological scrutiny.

2 New Racial Naturalism

2.1 Rebranding Naturalism: What Is New Racial Naturalism?

New Racial Naturalism (NRN) represents a strategic reformulation of the naturalist position on race. In contrast to Standard Racial Naturalism's reliance on biological essences or typological traits, NRN attempts to ground racial classification in population genetics—particularly in the clustering patterns revealed by genomic analysis. It discards essentialism while preserving the metaphysical claim that races correspond to biologically real, scientifically identifiable groups. Its ambition is to naturalize race without reducing it to an outdated 19th-century typology.

The most influential contemporary articulation of NRN comes from Quayshawn Spencer, whose work defends the view that race is real, biologically grounded, and structured by the U.S. Office of Management and Budget (OMB) racial classification system. According to Spencer (2012, 2014, 2020), these categories correspond to "genetically salient" continental-level populations that are revealed through analysis of multilocus genotype data using clustering tools like STRUCTURE. Spencer argues that when these algorithms are run with k $= 5$, they yield five major clusters that align with the five standard U.S. racial categories: African, European, East Asian, Native American, and Oceanian ancestry.

On Spencer's account, these clusters are not merely social constructions but objective biological entities—what he terms "Blumenbachian populations." These populations are said to be distinguishable using genetic data alone and, crucially, to possess sufficient biological integrity to count as natural kinds within the context of human population structure. The claim is not that these groups are essentialist in the old sense, but that they are stable, genetically informative, and taxonomically meaningful.

This is the core promise of NRN: that race can be biologically real without invoking essential traits or typological boundaries. Instead of looking for traits that all members of a race must share, NRN focuses on patterns of allele frequencies that cohere across populations and allow for group-level classification. These groupings are thought to reflect evolutionary history, geographic dispersion, and reproductive isolation over time—making them, in Spencer's view, biologically grounded and scientifically valid.

What distinguishes NRN from SRN, then, is its methodological sophistication and philosophical subtlety. It presents itself as a more defensible, empirically updated version of racial realism—one that aligns with contemporary population genetics and avoids the discredited assumptions of essentialist taxonomy. It does not deny that race is a social category, but it insists that this social role is underpinned by biological structure.

Yet as the sections that follow will demonstrate, NRN ultimately fails to deliver on this promise. Its biological categories are unstable, contingent on modelling assumptions, and shaped by political and bureaucratic frameworks rather than natural divisions. While NRN avoids the errors of SRN, it cannot escape the deeper problem: that the biological patterns it identifies do not justify the metaphysical claim that race is a natural kind. In practice, it collapses into a form of biologically informed constructivism, even as it denies that label.

2.2 Naturalizing Without Essences?

One of the defining features of New Racial Naturalism is its attempt to naturalize race while rejecting biological essentialism. Unlike Standard Racial Naturalism, NRN does not posit that members of a race share intrinsic traits that are necessary and sufficient for group membership. Instead, it draws on patterns of genetic similarity at the population level, hoping to avoid the metaphysical pitfalls that plagued earlier models. This move is theoretically attractive—but it raises a pressing question: can race be a natural kind without essences?

In philosophical terms, natural kinds are typically understood as categories that support inductive inference, possess explanatory power, and reflect objective divisions in nature. They need not be defined by necessary and sufficient conditions, but they must exhibit a level of internal coherence and causal unity. For race to qualify as a natural kind under NRN, its categories must not only be statistically robust, but also reflect biologically salient groupings that play an explanatory role in evolutionary or developmental processes.

NRN attempts to meet this standard by appealing to genetic clustering. When algorithms such as `STRUCTURE` are used to sort individuals into populations based on genotype data, they often identify clusters that correspond—approximately—to socially recognized racial categories. These clusters are interpreted by NRN advocates as revealing biologically real groupings, grounded in patterns of shared ancestry and historical reproductive isolation.

Yet this inference is unstable. First, the number and composition of clusters depend

on modelling parameters, such as the value of k (the number of clusters the algorithm is instructed to find), the choice of reference populations, and the specific genetic markers included. There is no unique, natural value of k, nor is there a principled reason to treat one clustering output as more ontologically valid than another. As Gannett (2010) and Kaplan and Winther (2014) have argued, the outputs of these models reflect statistical convenience and research design—not natural boundaries.

Second, the coherence of these clusters is often retrospective. That is, we interpret them as biologically meaningful because they align with existing racial categories, not because they emerge independently of them. This raises the spectre of circularity. If the categories we aim to vindicate are already socially constructed, and if our algorithms are tuned to detect patterns that resemble those categories, then the resulting clusters may simply reinforce pre-existing assumptions. The danger here is that we are not discovering biological races, but redescribing social races in the language of population genetics.

Third, even if certain clusters exhibit relative genetic differentiation, this does not establish that they are natural kinds. Differentiation is a matter of degree, not of type. It is well-established that human populations differ in allele frequencies across loci, often as a function of geographic distance and historical separation. But these differences are gradual and overlapping, not discrete. The absence of sharp boundaries renders racial categories, as defined by NRN, unfit for naturalistic taxonomy.

In this sense, NRN fails to provide the metaphysical scaffolding required to naturalize race. It offers clusters without kinds, structure without essence, and pattern without explanatory depth. By abandoning essentialism but retaining naturalism, NRN positions itself in an unstable conceptual space: one that borrows the language of modern biology but cannot support the ontological commitments it implies.

The result is a form of metaphysical minimalism: a view that attempts to say something about the reality of race while committing to as little as possible about what race really is. But natural kinds are not minimal commitments. They are robust conceptual tools, with explanatory, epistemic, and inferential utility. If NRN cannot demonstrate that its categories satisfy these conditions, then it fails to naturalize race in any philosophically meaningful sense.

2.3 Reference Populations and the Geography of Misclassification

A deeper problem for New Racial Naturalism lies in its reliance on reference populations to generate the clusters it treats as biologically real. Clustering algorithms such as `STRUCTURE`, which NRN heavily depends on, do not produce racial groupings in a conceptual vacuum. Rather, they produce outputs that are shaped by the initial selection of data: the individuals sampled, the markers chosen, and the populations designated as references. In short, the clusters that emerge are conditional on methodological framing—they reflect the structure of the dataset more than the structure of the species.

This dependence creates a geographic and historical bias in the construction of race-like clusters. Most large-scale genetic studies have disproportionately sampled individuals from specific regions—often reflecting political borders, funding priorities, or national health systems. These sampling decisions influence how population boundaries are drawn. A clustering analysis based on African, European, and East Asian reference samples may yield three broad

clusters. But include Oceanian, Middle Eastern, or Central Asian groups, and the picture becomes dramatically more complex. There is no stable "ground truth" to determine which clustering is correct.

More troublingly, the output of clustering algorithms can change dramatically depending on the value of k—the number of clusters the algorithm is told to find. At $k = 2$, clustering often reflects continental ancestry (e.g., African vs. non-African). At $k = 5$, the output may resemble U.S. Census categories. At $k = 12$ or higher, fine-grained regional populations emerge that bear little resemblance to racial typologies. The decision to stop at $k = 5$ and to treat that result as ontologically privileged is not dictated by biology. It is dictated by institutional convenience.[2]

Spencer's defence of NRN explicitly relies on this $k = 5$ result and ties it to U.S. Census categories. But this move exposes a deeper tension: if race is biologically real because it aligns with social classifications, then biology is being used to validate political categories—not to discover natural kinds. It is difficult to see how such an approach avoids circularity. The clustering reflects the census, the census reflects historical race concepts, and these concepts are then treated as vindicated by clustering.

Additionally, geographic ancestry is a moving target. Human migration, admixture, and globalisation have rendered any fixed mapping of genes to geography increasingly unstable. Individuals may carry genetic markers associated with multiple continental populations, yet identify with none of the standard racial categories. This is not merely a sociological point—it directly undermines the coherence of NRN's biological taxonomy. If individuals do not neatly fall into its clusters, then the categories fail as kinds.

Consider, too, the consequences of applying NRN across different national contexts. In Brazil, for instance, racial classification is based more on phenotypic expression than on presumed ancestry. In South Africa, categories such as "Coloured" represent complex histories of mixture that defy discrete clustering. NRN cannot explain or accommodate this variability without abandoning the claim that race is a natural kind. The dependence on U.S.-centric population references reveals that NRN is not globally applicable—it is a local ontology presented as a universal taxonomy.

In sum, NRN's appeal to reference populations and clustering algorithms fails to insulate it from the conceptual criticisms that defeated SRN. Its biological realism is conditional, path-dependent, and geographically partial. By mistaking artefacts of sampling for natural categories, NRN fails to provide a generalizable or principled account of race as a biological kind.

2.4 The Census Collapse: NRN as Policy, Not Taxonomy

One of the more revealing aspects of New Racial Naturalism is its reliance on the racial categories defined by the U.S. Office of Management and Budget (OMB). Spencer, in particular, grounds his account of race in what he calls "OMB race theory," arguing that the five official census categories—White, Black or African American, Asian, American Indian or Alaska Native, and Native Hawaiian or Other Pacific Islander—map onto biologically salient

[2]Indeed, in STRUCTURE analyses, varying k yields dramatically different outputs. The privileging of $k = 5$ reflects institutional convenience more than biological coherence—mirroring U.S. census categories.

human populations identified through genetic clustering at $k = 5$.

This move is rhetorically effective but philosophically untenable. Census categories are policy tools. They are shaped by historical contingencies, political negotiation, and pragmatic aims. They are not derived from biological investigation, nor were they ever intended to reflect taxonomic structure. To treat them as evidence of natural kinds is to invert the epistemic hierarchy: what ought to be an empirical discovery becomes a bureaucratic presupposition.

The circularity is obvious upon inspection. The census defines race in part based on social perception and self-identification. Genetic studies then map their findings onto these categories and claim to recover biologically real groupings. But if the categories themselves are socially constructed, and the clustering models are trained or interpreted in relation to them, then the output is merely echoing the input. What results is a feedback loop between policy and science, not an independent discovery of racial kinds.

This problem is not unique to NRN but is magnified by its methodological strategy. By explicitly tying race to U.S. governmental categories, NRN limits itself to a parochial framework that lacks cross-cultural applicability. The five OMB groups are not used globally, nor are they exhaustive of global human variation. Categories such as "Middle Eastern," "North African," or "South Asian" are excluded or ambiguously classified, despite comprising significant population groups. If race were a natural kind, its categories would not shift with bureaucratic regimes.

Moreover, the OMB system is itself unstable. Over time, its definitions and classifications have changed in response to demographic shifts and political pressure. For example, the inclusion of "Hispanic or Latino" as an ethnicity rather than a race, or the option to select "more than one race" in more recent censuses, reflects the evolving social landscape. This plasticity undermines any claim that these categories track stable biological kinds.

Even if one grants that the census categories correspond loosely to major continental ancestry groups, this concession does not support the metaphysical claim that they are natural kinds. The categories are at best heuristic proxies for genetic ancestry—not ontological divisions. Their utility in certain biomedical contexts, such as population stratification correction in genome-wide association studies, reflects pragmatic compromise, not taxonomic legitimacy.

To appeal to the census as a biological ontology is to mistake governance for metaphysics. It is to treat instruments of demographic classification as revelations of natural order. This is not only a category mistake—it is a philosophical failure to distinguish between epistemic convenience and ontological commitment.

The collapse of NRN into census policy exposes the thinness of its naturalism. If the best evidence for biological races is that they match U.S. bureaucratic categories, then the project of racial naturalism has abandoned the ambition of scientific universality. What remains is not a theory of natural kinds, but a biologically ornamented version of social sorting. In trying to save race as a biological kind, NRN merely repackages race as a policy artifact.

2.5 Is NRN Just Constructivism in Denial?

Despite its elaborate theoretical scaffolding, New Racial Naturalism ultimately fails to deliver on its central promise: that race is a biologically real kind, discoverable through the tools of

population genetics. As we have seen, the clusters it identifies are shaped by methodological assumptions, its categories derive from political institutions, and its metaphysical claims exceed what the science can support. What remains, beneath its technical apparatus, is a view that looks increasingly indistinguishable from biologically informed social constructivism.

To be clear, constructivism does not entail that race is arbitrary or unreal. It holds that racial categories are socially constructed, historically situated, and causally potent. They are real in the way that borders, money, or institutions are real: as collectively maintained structures that organize human life. Many constructivist accounts acknowledge that race draws on physical features and biological ancestry—but they deny that these associations reflect natural kinds. Race is real because it is constructed, enacted, and enforced—not because it is discovered.

NRN claims to reject this position. It insists that race is not merely socially imposed, but biologically grounded. Yet once essentialism is abandoned, and once explanatory depth is ceded to statistical patterning, it is difficult to see how NRN preserves the metaphysical distinction it claims. Its categories do not function as natural kinds in the taxonomic sense. They are neither explanatory nor predictive in ways that exceed their social reference points. They are read through the lens of historically constructed groupings, interpreted through culturally saturated classifications, and anchored in institutional regimes like the U.S. Census.

In practice, NRN uses biological data to trace the history of socially constructed groups. It shows that certain population clusters exist and that these clusters correlate with census categories, phenotypic features, and geographic ancestry. But this is precisely the kind of insight a constructivist can accept. Constructivists do not deny that social classifications leave biological traces—only that those traces amount to natural kinds. Admixture, migration, and reproductive sorting can produce biologically patterned populations without those populations becoming metaphysically significant taxa.

Indeed, much of NRN's empirical content is consistent with what Haslanger (2012), Gannett (2010), and Hardimon (2017) refer to as the material reality of race: the idea that socially constructed racial categories can have biological correlates without being biologically fundamental. This view allows for local biological differentiation—differences in allele frequency, ancestry, disease prevalence—without reifying race as a natural kind. It maintains a clear distinction between biological variation and racial taxonomy.

What distinguishes constructivism from NRN, then, is not the data, but the metaphysics. NRN insists on a realist interpretation of clustering outputs, while constructivists see those outputs as shaped by and responsive to social classification. The danger is that NRN confuses the map for the territory—mistaking correlation for kindhood, structure for essence, and clustering for taxonomy. Its realism is performative rather than philosophical: it asserts that race is natural while relying on a framework that denies the conditions under which kinds can be said to exist.

The upshot is this: NRN does not provide an alternative to constructivism. It is constructivism by another name—recast in genomic language, tied to policy categories, and backed by statistical modelling. It preserves the appearance of biological realism without its substance. In trying to defend race as a biological kind, it succeeds only in redescribing what constructivists have long understood: that race is biologically legible because it is socially inscribed.

If race is to be treated as a scientifically meaningful concept, it must be on terms that reflect its historical, social, and material construction—not through the repurposing of population genetics to validate bureaucratic categories. NRN, despite its sophistication, does not escape the gravitational pull of constructivism. It merely resists the label.

3 Race as a Social Construct: Toward Constructivist Realism

3.1 What It Means to Say Race Is Constructed

To say that race is socially constructed is not to claim that it is unreal, arbitrary, or fictitious. On the contrary, constructivist accounts of race hold that racial categories are real in a distinctively social way: they are produced and maintained through social practices, institutions, and discourses, and they exert causal force in the world. Understanding the construction of race requires conceptual clarity, not rhetorical flourish. We must ask: what kind of thing is a socially constructed kind, and what does it mean for such a kind to be real?

Constructivism is not a single thesis but a family of views. At minimum, constructivist accounts deny that race corresponds to a natural kind in biological taxonomy. Instead, they assert that racial categories are the product of historical, political, and cultural processes that impose meaning on human variation. This imposition is not merely linguistic or conceptual—it is institutional and material. Racial classifications shape legal systems, medical research, economic opportunity, and patterns of violence. They are sustained by social structures and enacted through lived experience.

To capture this complexity, recent social ontology has drawn distinctions among different modes of construction. Haslanger (2012) distinguishes between manifest and operative conceptions of race—the former capturing how race is overtly defined, the latter how it functions within systems of power. Ásta (2018) introduces the notion of "conferralist" social kinds, which come into being through acts of social recognition and classification. On her view, being racialized is not merely a matter of fitting a descriptive profile, but of being positioned within a socially sustained framework of norms and expectations.

This perspective aligns with Ian Hacking's (2006) notion of "looping kinds," in which classifications shape the behaviour, self-understanding, and institutional treatment of those they classify. Racial categories are not static labels; they are interactive constructs that evolve in response to how they are used.[3] Individuals become, resist, internalize, or transform their classifications, generating new forms of social meaning and material consequence. This reflexivity distinguishes social kinds from natural kinds, and it highlights the dynamism of race as a lived category.

To say that race is constructed, then, is to assert that it is historically emergent, socially maintained, and materially efficacious. It is not reducible to individual attitudes or discursive practices, nor is it invented ex nihilo. Racial classifications have genealogies: they are shaped

[3]Hacking's (2006) idea of "looping effects" helps explain how racial classifications shape the behaviour and institutions that then reinforce them.

by colonial histories, economic systems, legal regimes, and epistemic frameworks. Their meaning is sedimented across time and institutions, yet subject to revision and contestation. Race is real because it is collectively enforced and differentially experienced.

Importantly, social construction is not opposed to objectivity or causality. Constructed kinds can support explanation, generate predictions, and figure centrally in scientific inquiry. As Khalidi (2016) argues, social kinds may be more ontologically fragile than natural kinds, but they are no less real. Their metaphysical status derives from their causal roles and social scaffolding, not from their presence in nature independent of human practices.

What emerges is a robust framework for understanding race as a social kind: not a natural kind misidentified, but a social kind with biological footprints, historical inertia, and political salience. Constructivist realism—properly understood—does not diminish the reality of race. It clarifies its source, its structure, and its stakes.

3.2 The Reality of Social Kinds

If racial categories are not natural kinds, does that mean they lack ontological status? The answer, increasingly defended in contemporary metaphysics, is no. Social kinds—categories such as race, gender, money, citizenship, and institutions—are not less real than natural kinds. They are real in a different register: constituted by social practices, maintained through collective recognition, and embedded in material structures. Their reality is not grounded in nature independent of human activity, but in the patterned ways that human beings make and remake the world.

This view, often called "social kind realism," has gained significant traction in recent work by authors such as Amie Thomasson, Muhammad Ali Khalidi, and Ásta. On Thomasson's (2003, 2014) "easy ontology" account, social entities are ontologically legitimate if they meet minimal conditions of existence and application within our conceptual and linguistic practices. There is no need to postulate mind-independent essences to vindicate the reality of social kinds. If people recognize, interact with, and structure their lives around a category, and if that category performs a stable function in discourse and action, then it is real enough to warrant ontological recognition.

Khalidi (2016, 2020) offers a complementary account grounded in the causal-explanatory role of kinds. On his view, kinds are legitimate if they support successful explanations, enable prediction, and function coherently in scientific inquiry. Social kinds meet these standards when they are embedded in institutional contexts, generate regularities in behaviour or outcomes, and interact causally with other phenomena. Racial categories, for instance, help explain patterns in healthcare access, educational attainment, policing outcomes, and psychological stress. They are not mere labels but variables with demonstrable causal significance.

Importantly, the metaphysical dependence of social kinds on human activity does not entail ontological fragility. While it is true that social kinds can change or dissolve—consider the historical disappearance of certain social roles or categories—they often exhibit remarkable stability. Racial classifications, in particular, have persisted across centuries and institutional reforms. They are not invented anew in each interaction; they are inherited, taught, enforced, and contested. Their causal power is anchored in what Ásta (2018) calls "social conferral": the collective act of placing individuals within structures of significance

and expectation.

Moreover, social kinds frequently interact with material and biological systems. Racial classifications shape residential patterns, environmental exposures, healthcare biases, and legal outcomes—factors that, in turn, affect bodies and lives. As a result, race becomes biologically visible not because it is biologically fundamental, but because social processes inscribe themselves onto biology. This is especially evident in health disparities, where differential access and exposure manifest as physiological difference.

Thus, social kinds are not "merely" social. They are ontologically rich, causally embedded, and explanatorily indispensable. Denying their reality because they are socially constructed is akin to denying the reality of legal institutions or economic systems. What matters is not whether a kind is found in nature independent of us, but whether it plays a role in the world we inhabit and shape. On that score, racial categories qualify without hesitation.

In defending the reality of social kinds, we therefore defend the metaphysical plausibility of race as a socially constructed kind—one that is responsive to classification, causally efficacious, and institutionally entrenched. It is a kind that exists because we treat it as real, and in treating it as real, we make it so.

3.3 The Biological Surface of Social Construction

One of the most powerful objections to constructivist accounts of race is that racial categories track biological differences that appear to have medical or genetic significance. Studies show that disease prevalence, drug metabolism, and health outcomes vary across racial lines. These patterns seem to demand biological explanation—and appear, at first glance, to support some version of racial naturalism. But this inference mistakes the biological uptake of social inequality for the discovery of biological kinds.

Constructivist realism is not committed to denying that biology matters. It recognizes that socially constructed racial classifications can acquire biological expression. The mechanisms of this process are increasingly well understood. They include historical patterns of structural inequality, differential environmental exposure, and the embodiment of social stressors. As Dorothy Roberts (2011) and Jonathan Kahn (2013) argue, race in biomedicine often reflects the sedimentation of social disadvantage into the body, not the mapping of biological essence onto social identity.

Health disparities provide a clear case in point. In the United States, Black patients experience higher rates of hypertension, diabetes, stroke, and maternal mortality. While these differences are often racialized in clinical discourse, they do not reflect genetic groupings so much as the cumulative effects of social stress, environmental racism, inadequate care, and systemic bias. The biological variation is real—but it is socially produced.

This process is now being illuminated by epigenetics: the study of how gene expression is regulated by environmental and social factors. Epigenetic mechanisms—such as DNA methylation or histone modification—respond to stress, diet, pollution, and trauma, many of which are unequally distributed along racial lines. In this light, biological difference becomes a downstream consequence of lived inequality. Race becomes legible on the body not because it originates there, but because society marks bodies through exposure, access, and risk.

This is not a new insight.[4] W.E.B. Du Bois, writing in the early 20th century, argued that racial disparities in health were not due to innate biological inferiority but to "the vast structural differences in housing, education, income, and health provision." What modern science has done is clarify the pathways by which social inequality becomes biologically encoded. It has given empirical traction to Du Bois's philosophical claim.

Importantly, none of this undermines the constructivist view. In fact, it confirms it. The biological relevance of race is a function of social sorting. Racial categories do predictive and explanatory work in medicine and public health not because they are natural kinds, but because they track historically entrenched pathways of differential treatment, environmental exposure, and institutional neglect. Their epistemic utility is derivative, not fundamental.

Constructivist realism can accommodate this fact by distinguishing between biological significance and biological essentialism. It allows that racial categories may correlate with biological variation in specific contexts, without concluding that these categories carve nature at its joints. Race matters in health because racism matters in health—not because race is a biologically grounded kind.

This perspective reframes the role of race in science. Rather than treating it as a taxonomic category in need of biological vindication, we can treat it as a sociohistorical variable with material consequences. This allows researchers to investigate the causes of racialized health disparities without reifying race itself. It also aligns with the broader goal of constructivist realism: to explain how socially constructed categories can shape, and be shaped by, the world.

3.4 Against Eliminativism

Given the conceptual and empirical difficulties surrounding racial classification, some philosophers have proposed that we abandon the concept of race altogether. This position, known as racial eliminativism, argues that because race does not refer to a natural kind, and because its continued use perpetuates harm, we should discard it from scientific, philosophical, and political discourse. While well-intentioned, this response is both conceptually premature and politically misguided. Eliminativism underestimates the causal and normative role race continues to play in the world.

Eliminativist arguments typically follow one of two lines. The first is epistemic: since race fails to correspond to any robust biological kind, it is scientifically illegitimate and should be eliminated as a category of inquiry. This view is most closely associated with Kwame Anthony Appiah (1992) and Naomi Zack (2002), both of whom argue that continued reference to race reinforces false beliefs about human biology. The second line is normative: given that racial categorization has historically supported oppression, hierarchy, and violence, we ought to abandon the very terms that sustain these frameworks.

Both positions rest on a mistaken assumption: that eliminating the concept of race will eliminate the consequences of racialization. But as Haslanger (2012) and Hardimon (2017) argue, this conclusion fails to grasp the structure of social kinds. Racial categories, even if not biologically grounded, are real in virtue of their social construction, institutional en-

[4]Du Bois's sociological analysis of race and health disparities—especially in his 1906 paper *The Health and Physique of the Negro American*—predated contemporary epigenetics by nearly a century.

trenchment, and material impact. They structure access to resources, shape lived experience, and guide institutional behaviour. To erase the concept without dismantling the structures it names is to misidentify the source of the problem.

Moreover, eliminativism risks rendering invisible the very injustices it seeks to overcome. If racial disparities in policing, health, housing, and education are the result of racialized structures, then analysis and redress require a conceptual framework that tracks those structures. Without race as an analytic category, it becomes difficult to name, challenge, or rectify patterns of systemic inequality. Eliminativism offers moral purity at the cost of political clarity.

Constructivist realism provides a more coherent alternative. It acknowledges that race is not a natural kind while affirming its social reality. It maintains that racial categories are constructed and revisable, but not arbitrary or dispensable. It treats race as a tool for understanding the world we have made—a world in which classification itself has become causal. From this perspective, the goal is not to eliminate the concept of race, but to reinterpret it in a way that is philosophically responsible and politically effective.

This approach also better reflects the lived epistemology of racialized subjects. As Charles Mills (1998) has argued, those who experience the consequences of race cannot afford to treat it as an illusion. To do so is to disavow the epistemic authority of those most affected by racial categorization. Constructivist realism affirms that race is both a system of classification and a structure of experience—one that is constructed, real, and actionable.

Finally, eliminativism struggles to account for the productive uses of racial identity. In many contexts, race is not simply imposed but claimed, reinterpreted, and politicized. From anti-colonial resistance movements to contemporary struggles for justice, racial identity can serve as a source of solidarity and critique. Eliminating the concept risks eliminating a crucial axis of collective agency. Race, on this view, is not only an obstacle to be overcome, but also a resource to be reappropriated.

In sum, the eliminativist impulse fails to appreciate the ontological complexity and political utility of race. Constructivist realism, by contrast, offers a framework that is empirically grounded, philosophically coherent, and normatively engaged. It rejects the false choice between realism and elimination, and insists that understanding race requires confronting it—not erasing it.

3.5 Constructivist Realism and the Ontology of Resistance

Constructivist realism offers not only a metaphysical diagnosis of racial classification, but a philosophical framework capable of supporting critique, resistance, and political transformation. It affirms that race is real—not because it reflects natural divisions within the human species, but because it structures the world we inhabit. It is real in its consequences, institutions, and embodiment. And it is real in the possibilities it opens for reclaiming agency, identity, and justice.

The strength of constructivist realism lies in its dual orientation. It is at once ontological and normative. It takes seriously the causal and epistemic roles racial categories play, while remaining critical of the forces that sustain them. It recognizes that race is constructed through systems of power and domination, but does not conclude that its meaning is exhausted by oppression. Constructivist realism allows us to understand race as a product of

injustice without committing to its permanence or inevitability.

This orientation distinguishes it from both racial naturalism and eliminativism. Where the former seeks to vindicate race by locating it in biology, and the latter seeks to reject race by denying its reality, constructivist realism locates race in the social world—with all the complexity, contingency, and mutability that entails. Race is neither an illusion nor a natural kind. It is a social kind with material inertia and ethical gravity.

Importantly, this view does not commit us to a static ontology. As Ásta (2018) and Haslanger (2012) have emphasized, social kinds can be contested, revised, and reconfigured. The ontology of race is not fixed. It is responsive to collective action, discursive intervention, and political struggle. This makes race not only a category of classification, but a terrain of contestation. Understanding its construction is a precondition for its transformation.

In this light, constructivist realism provides the resources for an ontology of resistance. It enables us to name and critique racial hierarchies, without falling into the trap of biological reification. It grounds anti-racist politics in a sober metaphysical understanding of how race operates—how it is made, maintained, and potentially unmade. As such, it bridges the gap between descriptive analysis and normative commitment. It does not ask us to believe in race as nature, nor to pretend that race is nothing. It asks us to understand race as a historically situated system of classification with real effects and real stakes.

This framework also preserves the space for strategic deployment of racial categories. As Sally Haslanger has argued, acknowledging the constructedness of race does not preclude its use in emancipatory projects. On the contrary, understanding race as socially constructed enables us to interrogate the terms of its construction—and to reconstruct it in more just and inclusive ways. Constructivist realism allows for a politically conscious engagement with race that is both conceptually precise and ethically oriented.

Ultimately, the goal is not to redeem race as a timeless essence, nor to erase it as a historical mistake, but to understand it as a contingent structure of human sociality—one whose continued existence depends on our collective practices and whose transformation is within the reach of critique and action. Constructivist realism gives us the tools to think race otherwise.

It is not a retreat from ontology. It is an insistence that our ontologies be responsive to the world as it is—and as it might become.

4 Race as a Regulatory Kind: A Metaphysical Elaboration

The constructivist realist position developed above offers a compelling account of race as a socially constructed category with material consequences and institutional persistence. However, an outstanding question remains: how should we understand the kind-like stability that racial classifications exhibit in contexts such as forensic genetics, biomedical research, and policy design?

To address this, I propose that race is best understood as a *regulatory kind*: a socially constructed kind whose epistemic tractability and classificatory stability are sustained through recursive institutional processes. Whereas ideologically saturated kinds (Haslanger), social

role kinds (Mallon), and interactive kinds (Hacking) each account for the formation and social persistence of race, they do not fully isolate the mechanisms by which race can simulate the epistemic profile of a biological kind.

> **Regulatory Kind Schema.** For any socially constructed kind K, K is a *regulatory kind* if and only if:
>
> 1. There exists a set of institutional structures I (e.g., census protocols, forensic databases, biomedical coding systems) such that:
>   (a) I takes socially constructed input classifications S (e.g., administrative or census categories),
>   (b) processes them using classification procedures P (e.g., clustering algorithms, typological schema), and
>   (c) generates output data O that are then re-used to justify the continued application or refinement of K.
> 2. This produces a feedback relation R over time such that:
>
> $$R : S \xrightarrow{P} O \rightarrow S'$$
>
>   where S' reinforces the institutional uptake and apparent naturalness of K.
> 3. The perceived kindhood of K is therefore not grounded in any natural property N, but rather in the stability produced by R under the operation of I.

This schema captures the ontological structure underlying practices that render race empirically robust in certain contexts, despite its lack of biological grounding. The account does not reject the insights of existing constructivist models but rather specifies a distinct subtype of socially constructed kind—one whose stability is epistemically engineered rather than biologically discovered.

This also clarifies the failure of New Racial Naturalism. NRN's use of clustering algorithms and population partitions appears to uncover biological groupings, but these outputs often reflect prior social classifications embedded in the input data. When these outputs are then used to inform policy, research, or social categorisation, they reinforce the initial assumptions, completing the feedback loop. On the present account, such outputs do not uncover natural kinds; rather, they instantiate the recursive logic of regulatory kindhood.

Understanding race as a regulatory kind allows us to preserve the metaphysical commitments of constructivist realism while accounting for the empirical stability of racial categories in scientific and administrative domains. It explains why race appears stable and kind-like in classificatory practices without relying on essentialist assumptions, and highlights how epistemic authority itself can contribute to the construction of socially durable categories.

4.1 The Simulation of Kindhood: A Metaphysical Analysis

Regulatory kinds exhibit the appearance of natural kindhood across scientific and administrative domains. They seem to be stable, predictive, and projectible—hallmarks of what

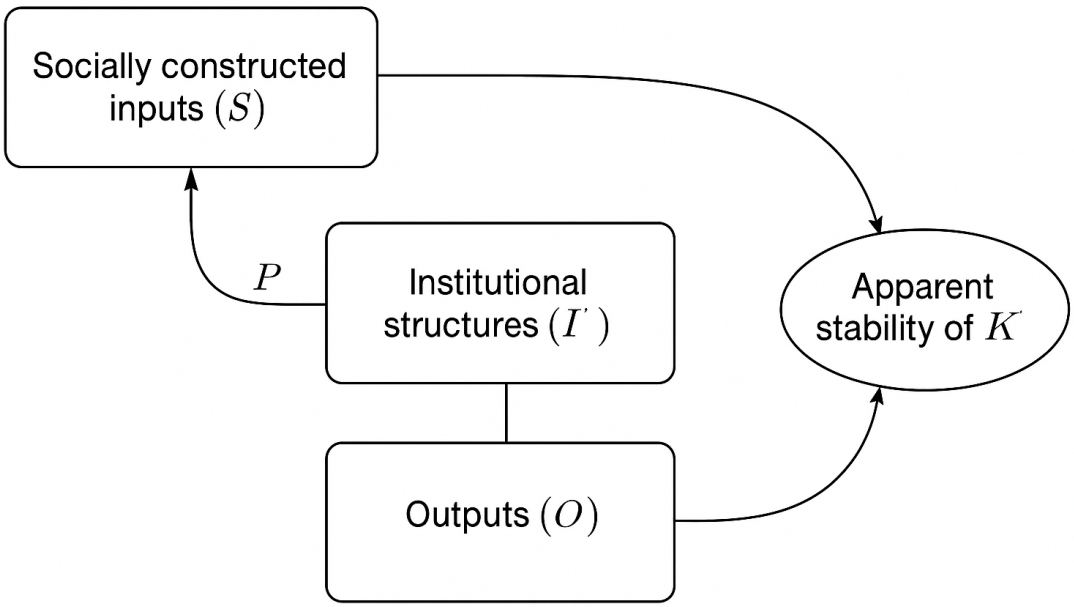

Figure 1: The recursive classification mechanism sustaining regulatory kindhood. Socially constructed inputs (S) are processed by institutional structures (I) via classification procedures (P), generating outputs (O) that recursively reinforce the apparent stability of K.

philosophers take to be natural kinds. Yet this appearance is deceptive. As outlined in the schema introduced in Section 4.1, regulatory kinds generate this apparent stability not through essential properties or natural kindsmanship, but through recursive institutional feedback mechanisms. Their kindhood is not discovered; it is enacted.

Definition (Kindhood Simulation). Let K be a socially constructed kind. K *simulates natural kindhood* if and only if:

(S1) Epistemic Stability: K supports inductive generalizations across domains (e.g., in medicine, forensics, education).

(S2) Cross-System Coherence: K is used consistently across multiple institutional and epistemic systems.

(S3) Recursive Classification: K's apparent reliability is produced by a feedback mechanism wherein output data O are recursively used to justify and refine the classification inputs S.

This simulation structure directly maps onto the recursive feedback loop R specified in Section 4.1:

$$R : S \xrightarrow{P} O \to S'$$

where S denotes socially constructed input classifications, P denotes classification procedures, O denotes output data, and S' denotes revised or reinforced input categories. As this loop stabilizes under the operation of institutional structures I, it generates the illusion that K reflects an underlying natural kind. This is the core mechanism of regulatory kindhood: epistemic tractability manufactured through recursive enforcement.

Bridge Principle (Simulation-to-Kindhood Misidentification).
If a kind K satisfies (S1–S3) and its institutional outputs O are interpreted as evidence of biological or natural kind status, then there is a risk of *kindhood misidentification*: the error of treating simulated stability as metaphysical depth.

Corollary (NRN Failure Condition).
If New Racial Naturalism (NRN) treats clustering outputs O—which are produced by procedures P applied to socially constructed inputs S—as indicators of a natural property N, then NRN commits kindhood misidentification unless N can be shown to exist independently of I, P, and R.

This analysis clarifies the metaphysical error underlying NRN. Clustering algorithms (e.g., `STRUCTURE`) are sensitive to assumptions about k, population selection, and reference inputs—all of which are shaped by regulatory classification systems. When outputs coincide with racial categories, this reflects the simulation of kindhood, not its discovery.

The recursive mechanism R generates outputs O that appear biologically meaningful, but which are in fact epistemically downstream of socially constructed inputs S. These outputs then justify further applications of the same classificatory regime, reinforcing the illusion of objectivity. Over time, this feedback loop stabilizes a kind that behaves like a natural kind while lacking any metaphysical grounding in nature.

> *If all K-classifications are processed through R and appear stable under I, then K will display the epistemic virtues of natural kindhood even if no natural property N grounds K.*

This helps explain why racial categories can perform well in biomedical and forensic contexts. Their success is not accidental—but neither is it natural. It is the result of recursive design. The scientific utility of race is engineered through the regulatory architecture described in Section 4.1.

By diagnosing this structure, we do not deny the empirical salience of race. Rather, we relocate its source: from biology to bureaucracy, from nature to recursive institutional design. Regulatory kinds simulate kindhood by rerouting epistemic authority through classification regimes. To mistake this simulation for essence is to misread infrastructure as ontology.

4.2 Distinctiveness from Other Social Kind Models

The concept of regulatory kinds intersects with, but crucially diverges from, several major models in the metaphysics of social kinds. This section clarifies these distinctions, demon-

strating that regulatory kinds are not merely a subspecies of existing theories, but a structurally distinct class defined by recursive institutional enforcement and epistemic durability.

1. Interactive Kinds (Hacking)
Hacking (2006) introduced the idea of interactive kinds—categories that affect the people classified under them, who in turn modify the category itself. These kinds "loop" due to the feedback between classification and self-understanding.

Contrast: Regulatory kinds extend beyond looping effects by embedding kind formation in institutional infrastructures. They need not depend on the subject's self-understanding. What matters is the operation of recursive classification via institutional systems.

Formal Divergence: Let K be a kind. If K is maintained through individual or group uptake of classification (e.g., internalization, resistance), then K is an interactive kind. If K is stabilized through recursive institutional feedback regardless of uptake, then K is a regulatory kind.

2. Ideologically Saturated Kinds (Haslanger)
Haslanger (2012) describes certain social kinds—especially gender and race—as constructed through ideology and social meaning. Their persistence is explained by normative expectations and power-laden interpretive frames.

Contrast: While regulatory kinds are often ideologically saturated, their ontological stability derives not from meaning or normativity alone, but from recursive bureaucratic implementation. They may persist even when ideological content is weak, contested, or denied.

Formal Divergence: Let K be a kind. If K's persistence depends primarily on shared ideology or social scripts, then K is ideologically saturated. If K is stabilized via institutional procedures that classify, record, and reapply K, then K is regulatory—even if ideological content is minimal.

3. Social Role Kinds (Mallon)
Mallon (2016) defines social role kinds as kinds constructed through normative roles and expectations assigned within specific institutions. These are position-based, involving behavioral scripts and social structures.

Contrast: Regulatory kinds are not reducible to roles. A person may be assigned to a racial classification without occupying any coherent social role. The kindhood emerges from data processing, policy recirculation, and epistemic uptake—not from norm-guided action.

Formal Divergence: Let K be a kind. If K depends on an individual fulfilling a normative role in a relational structure, it is a social role kind. If K requires only classification and recursive enforcement across systems, then it is a regulatory kind—even if no role is enacted.

4. Explanatory Kinds (Khalidi)
Khalidi (2020) argues that kinds are legitimated by their explanatory and predictive success within scientific theories and practices.

Contrast: Regulatory kinds may yield high explanatory utility, but this is often the result of

classification infrastructure rather than ontological grounding. Their success is engineered—produced by systems that sort and analyze populations recursively.

Formal Divergence: Let K be a kind. If K's explanatory utility emerges independently of recursive classification feedback (e.g., natural phenomena), then K may be an explanatory kind. If K's explanatory success derives from institutional generation and reinforcement of data about K, then K is regulatory and its utility is infrastructurally contingent.

Summary Table: Comparison of Social Kind Models

Kind Model	Stabilizing Mechanism
Interactive (Hacking)	Self-understanding and social uptake loop back into the classification
Ideologically Saturated (Haslanger)	Normative expectations and shared ideology enforce category meaning
Social Role (Mallon)	Individuals occupy normative positions within relational institutional structures
Explanatory (Khalidi)	Scientific explanatory success justifies kindhood
Regulatory (Brewer)	Recursive institutional enforcement stabilizes kindhood via data, policy, and classification feedback

This contrastive mapping clarifies the conceptual terrain. Regulatory kinds overlap with these models, but they isolate a distinct mechanism of kind formation: epistemic enforcement through recursive institutional infrastructure. They do not merely persist due to social roles, meaning, uptake, or utility. They persist because bureaucratic systems depend on their classification to function, and in functioning, reproduce the classification.

This makes regulatory kinds particularly resilient to metaphysical critique. One cannot dissolve them merely by questioning their naturalness or ideological legitimacy. Their reality is a function of infrastructural entrenchment. Understanding this distinguishes race not just as a socially constructed kind, but as a kind that simulates naturalness because it is produced and maintained through mechanisms of regulation.

4.3 Regulatory Kinds Beyond Race? A Philosophical Coda

While this paper has focused on race as the central case, the framework of regulatory kinds is not meant to apply exclusively or exhaustively to racial classification. Rather, it offers a metaphysical model for understanding how certain socially constructed kinds can acquire classificatory stability, empirical traction, and institutional durability through recursive epistemic mechanisms.

This model may have broader application. Any social kind that satisfies the core conditions outlined in Section 4.1—recursive feedback via institutional classification, cross-domain uptake, and epistemic reuse—may, in principle, instantiate regulatory kindhood. Whether a particular kind qualifies is a substantive, domain-specific question that should be determined by empirical analysis and normative caution.

The aim here is not to universalize the model, but to offer a conceptual architecture that may be of use elsewhere. If scholars working in fields such as disability studies, psychiatric classification, immigration policy, or educational stratification find that their domains exhibit similar recursive patterns, the concept of regulatory kinds may serve as a helpful analytic tool. But that uptake must be earned, not assumed.

Race remains the paradigmatic case. It exemplifies how a socially constructed kind can behave with the epistemic rigor of a natural kind—despite lacking natural essence—through recursive enforcement by census categories, forensic protocols, health disparities research, and administrative records. Understanding race as a regulatory kind clarifies both its empirical grip and its metaphysical contingency.

This framework does not seek to displace other models of social kindhood, but to supplement them with a tool for analyzing classification regimes that simulate kindhood through infrastructural recursion. Whether others adopt it remains to be seen. The task of social metaphysics is not only to categorize, but to offer concepts that help us see how classification itself becomes ontological.

Conclusion: Ontology Without Essence

This paper has examined the prospects for racial naturalism in both its classical and revised forms. Standard Racial Naturalism, grounded in essentialist assumptions about biological difference, fails to meet the empirical and metaphysical standards required of a scientifically coherent taxonomy. New Racial Naturalism, while more subtle in its appeal to population structure and statistical clustering, ultimately collapses into a metaphysically unstable position. Its apparent empirical success is a simulation of kindhood, generated by recursive classification systems that reprocess socially constructed inputs into scientifically legible outputs.

In rejecting racial naturalism, however, this paper does not adopt an eliminativist stance. Race is real—not as a natural kind, but as a socially constructed, materially consequential, and institutionally enforced kind. It is precisely this recursive enforcement that gives rise to its empirical durability and epistemic tractability. Race is not biologically grounded, but it is bureaucratically stabilized. Its kindhood emerges not from nature, but from regulation.

To capture this ontological structure, I have introduced the concept of a *regulatory kind*: a socially constructed kind whose classificatory stability is maintained through recursive institutional feedback. Regulatory kinds simulate the epistemic behavior of natural kinds without being grounded in natural essences. They achieve projectibility, coherence, and cross-domain stability through data architecture, administrative protocols, and infrastructural recursion.

Race exemplifies this form of kindhood. Its classification is enforced through census categories, medical coding, forensic practices, and social policy. These systems do not merely reflect race—they reproduce it. They generate the outputs that appear to confirm race's kindhood, closing the epistemic loop. In understanding race as a regulatory kind, we make sense of its empirical success without conceding to naturalism, and we resist eliminativism without relying on essentialism.

The framework developed here is not proposed as a universal theory of social kinds, but as

a conceptual tool. If other domains—psychiatric diagnosis, immigration status, educational classification—exhibit similar recursive architectures, the model of regulatory kinds may prove useful. Whether or not it travels, its explanatory power in the case of race is, I hope, clear.

Ontology, in this account, is not a reflection of nature's joints, but an analysis of classification's machinery. Race is real—but its reality lies not in biology. It lies in systems that sort, code, and reclassify; in histories sedimented into policy; in bodies shaped by exposure and access. To understand race is to understand these systems—and to change them, we must begin by naming the kind of kind that race has become.

References

Appiah, K. A. (1992). *In My Father's House: Africa in the Philosophy of Culture*. Oxford University Press.

Ásta. (2018). *Categories We Live By: The Construction of Sex, Gender, Race, and Other Social Categories*. Oxford University Press.

Barnes, E. (2020). Metaphysically engineering better social kinds. In M. Green & D. Plunkett (Eds.), *Conceptual Engineering and Conceptual Ethics* (pp. 351–374). Oxford University Press.

Bowker, G. C., & Star, S. L. (2000). *Sorting Things Out: Classification and Its Consequences*. MIT Press.

Edwards, A. W. F. (2003). Human genetic diversity: Lewontin's fallacy. *BioEssays*, 25(8), 798–801.

Gannett, L. (2010). Questions asked and unasked: How by doing genetic research we talk about race. In B. Koenig, S. Lee, & S. Richardson (Eds.), *Revisiting Race in a Genomic Age* (pp. 42–65). Rutgers University Press.

Hacking, I. (1991). A tradition of natural kinds. *Philosophical Studies*, 61(1–2), 109–126.

Hacking, I. (2006). *The Social Construction of What?* Harvard University Press.

Hardimon, M. O. (2017). *Rethinking Race: The Case for Deflationary Realism*. Harvard University Press.

Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford University Press.

Hochman, A. (2013). Against the new racial naturalism. *Journal of Philosophy*, 110(6), 331–351.

Kahn, J. (2013). *Race in a Bottle: The Story of BiDil and Racialized Medicine in a Post-Genomic Age*. Columbia University Press.

Khalidi, M. A. (2016). Social kinds and psychological kinds. *Mind & Language*, 31(5), 591–613.

Khalidi, M. A. (2020). *Kinds: A Philosophical Study of Classification*. Oxford University Press.

Lewontin, R. C. (1972). The apportionment of human diversity. *Evolutionary Biology*, 6, 381–398.

Mayr, E. (1963). *Animal Species and Evolution*. Harvard University Press.

Mills, C. W. (1998). *Blackness Visible: Essays on Philosophy and Race*. Cornell University Press.

Okasha, S. (2002). *Philosophy of Science: A Very Short Introduction*. Oxford University Press.

Roberts, D. (2011). *Fatal Invention: How Science, Politics, and Big Business Re-create Race in the Twenty-First Century*. The New Press.

Sankar, P. (2012). Forensic DNA phenotyping: Continuity and new challenges in the history of race, genetics, and policing. *New Genetics and Society*, 31(2), 195–207.

Sober, E. (1994). *Conceptual Issues in Evolutionary Biology* (2nd ed.). MIT Press.

Spencer, Q. (2012). What 'biological racial realism' should mean. *Philosophical Studies*, 159(2), 181–204.

Spencer, Q. (2014). A radical solution to the race problem. *Philosophy of Science*, 81(5), 1025–1038.

Spencer, Q. (2020). The race concept: A defense. *Studies in History and Philosophy of Science Part C*, 82, 101265.

Thomasson, A. L. (2003). Realism and human kinds. *Philosophy and Phenomenological Research*, 67(3), 580–609.

Thomasson, A. L. (2014). Ontology made easy. *Oxford Studies in Metaphysics*, 9, 78–105.

Zack, N. (2002). *Philosophy of Science and Race*. Routledge.