

Regulatory Kinds: A Metaphysical Framework for Epistemically Stabilized Social Classification

Mark A. Brewer

May 2025

Abstract

This paper introduces the concept of *regulatory kinds* — socially constructed classifications that come to function epistemically like natural kinds through recursive uptake across institutional domains. These kinds do not reflect causal unity or semantic precision, but they acquire stability, portability, and predictive utility by being embedded in the inferential routines of medicine, law, policy, and science. I develop the notion of *simulated kindhood* to explain how such classifications support explanation and coordination despite lacking metaphysical integrity. Race serves as the central case: a contested and heterogeneous category that nonetheless endures as a diagnostic tool, a policy metric, and a risk factor. By treating race as a regulatory kind, the paper reframes classificatory persistence as an institutional phenomenon, rather than a cognitive or conceptual error. The account challenges traditional views of kindhood, highlights the epistemic logic of infrastructural classification, and raises ethical concerns about the reification of simulated categories.

Keywords: social kinds; regulatory kinds; race; constructivist realism; institutional classification; philosophy of science; epistemic stability; metaphysics of categories

Introduction

Some social classifications behave more like scientific categories than social constructs are expected to. Categories such as race, disability, or criminality—though historically contingent and socially constructed—often exhibit a surprising degree of epistemic stability and predictive utility across domains such as policy, medicine, and law. This poses a metaphysical puzzle: how do socially constructed categories come to function “as if” they were natural kinds?

This paper introduces the concept of a *regulatory kind* to address this puzzle. Regulatory kinds are classifications that gain kind-like properties—such as explanatory traction, inferential power, and cross-contextual stability—not by virtue of any intrinsic natural essence, but through recursive processes of institutional uptake and enforcement. They are governed, reinforced, and rendered durable by the very systems that deploy them. This institutional recursion simulates the epistemic behavior of natural kinds while remaining grounded in historically and socially produced classifications.

The framework developed here builds on constructivist realism, the view that social constructs can have material consequences and causal power without being natural kinds in the traditional sense. But it goes further by offering a formal metaphysical model of how certain social constructs come to behave *as if* they were kinds. This approach differs from familiar accounts of social kinds—such as Ian Hacking’s looping kinds, Sally Haslanger’s ideologically saturated kinds, and Richard Boyd’s naturalized kinds—by focusing not on identity feedback or ideological embedding alone, but on the recursive stabilization of categories across multiple institutional settings.

Importantly, this account does not depend on a particular semantic theory of race or other contested terms. While recent work in the philosophy of race has produced divergent accounts of what “race” means in contemporary discourse—whether descriptivist (Hardimon 2017), referentialist (Spencer 2014), or pluralist (Spencer 2019)—this paper remains neutral on such debates. Its concern is not what these categories mean, but how they *behave* once institutionalized. The regulatory kind framework is intended to apply regardless of the semantic content of the classifications it models.

The argument proceeds in five sections. Section 1 examines the ontology of social kinds, tracing the limitations of existing metaphysical models — including interactive kinds, ideologically saturated kinds, and natural kind pluralism — in accounting for classificatory persistence and epistemic function. Section 2 introduces the framework of regulatory kinds and develops the concept of simulated kindhood: the idea that certain socially constructed classifications come to function epistemically like kinds through recursive uptake and institutional embedding. Section 3 applies this framework to the case of race, demonstrating how a semantically unstable and politically contested category acquires epistemic traction across domains such as medicine, public policy, and genomics. Section 4 explores the broader implications of regulatory kinds: their institutional reach beyond race, their risks of reification, and the ethical demands they impose on classification practices. Section 5 concludes by reflecting on the philosophical and political significance of simulated kindhood in institutional life.

The aim of this paper is not to defend the legitimacy of any specific social classification, nor to obscure the political work such classifications perform. Rather, it is to develop

a metaphysical account of how some classifications—regardless of their origins—come to function epistemically like kinds. Regulatory kinds offer a way of understanding the material and inferential power of socially constructed categories without lapsing into essentialism or conceptual inflation.

1. Constructivist Realism and the Ontology of Social Kinds

A central challenge in contemporary metaphysics of science is to account for the reality of classifications that are socially constructed, epistemically useful, and materially consequential. Nowhere is this challenge more apparent than in the case of social kinds: classifications such as race, gender, and disability that lack natural essences but appear to exhibit kind-like behavior across institutional and scientific domains. What distinguishes these classifications from arbitrary labels? What makes them stable enough to support explanation, prediction, and policy action—even when their ontological foundations are widely contested?

This question is at the heart of *constructivist realism*, an approach that has gained increasing traction in both philosophy of science and social ontology. Constructivist realists argue that certain social kinds are real not because they correspond to intrinsic natural properties, but because they function causally and epistemically in our best institutional and scientific practices. On this view, social kinds can be indispensable to the structure of inquiry and social coordination, even if their boundaries are porous, historically contingent, and politically fraught.

Sally Haslanger (2012) famously describes race and gender as *ideologically saturated* kinds: categories produced and sustained by systems of power that materially shape lives through their classification. On this model, classification is a vehicle for social stratification, and kindhood is a function of institutional entrenchment and norm enforcement—not biological or psychological universals. Haslanger’s view is explicitly critical: it does not seek to legitimize race and gender as scientific kinds, but to explain how they operate as socially constructed mechanisms of oppression.

A different but complementary strand of the literature focuses on the epistemic function of kinds. Muhammad Ali Khalidi (2013) argues that natural and social kinds should be treated symmetrically: both can be characterized by their explanatory role in science. He proposes a functionalist model of kindhood: a category is a kind when it plays a sufficiently central role in inductive and explanatory practices. In later work, Khalidi (2016) defends the legitimacy of social kinds even when their referents vary across contexts, provided they meet criteria for explanatory reliability. On this view, kindhood is not undermined by conceptual contestation, provided the classification supports successful scientific practices.¹

¹For a closely related model, see Richard Boyd’s (1999) account of natural kinds as “homeostatic property

Ian Hacking (1999), meanwhile, emphasizes the dynamic relation between classification and the classified. His “looping kinds” model describes human kinds whose members change their behavior in response to how they are categorized. The classification thus shapes the phenomenon it purports to describe, introducing a feedback loop unique to the human sciences. This view has been widely influential in sociology, psychology, and the philosophy of psychiatry, where diagnostic categories are often shaped by public perception, clinical practice, and patient response.

Each of these approaches captures a crucial aspect of social kindhood. Haslanger foregrounds the normative force of classification; Khalidi and Boyd focus on explanatory robustness; Hacking emphasizes reflexivity. But these models tend to focus either on *intra-personal effects* (e.g., identity or behavior) or on *disciplinary knowledge systems* (e.g., science or medicine). What is often missing is an account of how social kinds operate *across* institutional domains—how they become stabilized and portable even when they are neither naturally grounded nor ideologically consistent.

Consider the classification of race in contemporary institutions. In the United States, racial categories appear in demographic surveys, hospital intake forms, educational policy, legal documents, and employment records. These categories are neither biologically coherent nor semantically uniform, yet they are treated as stable across contexts. The same classification—“Black,” “White,” “Asian”—guides decisions in healthcare algorithms, forensic databases, census policy, and risk modeling. What explains the stability and portability of these classifications across epistemically diverse systems?

This is the gap the present account seeks to fill. The concept of a *regulatory kind* is introduced to model socially constructed classifications that simulate kindhood across institutional contexts. A regulatory kind is not grounded in intrinsic properties (as in natural kinds), nor in looping behavioral feedback (as in Hacking’s model), nor solely in ideological embedding (as in Haslanger’s). Rather, it is stabilized by *recursive institutional uptake*: systems repeatedly deploy the classification as if it were stable, coordinating decisions, structuring data collection, and reinforcing the category’s legitimacy through practice.

This distinguishes regulatory kinds from familiar alternatives. Unlike Boyd’s natural kinds, which are stabilized by causal mechanisms that underwrite inductive generalization (e.g., the tree of life in biology), regulatory kinds lack any underlying natural or mechanistic basis. Their stability is not due to natural homeostasis, but to institutional recursivity. This distinction is not merely ontological but epistemological: Boyd’s kinds are assumed to support reliable inference because they track real causal patterns in the world. Regulatory kinds, by contrast, achieve apparent inferential power through socially sustained simulation.

clusters” stabilized by background causal mechanisms. While Boyd’s view was developed with biological taxa in mind, he explicitly allows for social kinds that satisfy the same explanatory role, such as economic categories or legal identities.

Their epistemic reliability is contingent, context-dependent, and often politically charged.

Unlike Haslanger’s ideological kinds, they are not solely vehicles of normative power. And unlike Hacking’s looping kinds, they do not depend on individual identification or behavioral change. What they share is a pattern of *epistemic simulation*. Their kindhood is not discovered but enacted—sustained not by natural law, but by coordinated classification practices. Their explanatory and predictive utility is a product of recursive classification, not natural essence.

Regulatory kinds are not defined by personal identification, feedback responsiveness, or normative power alone. Nor are they merely any classification that recurs across institutions. What distinguishes them is a specific structure: the category is treated by multiple systems as if it supports reliable generalizations, even when its ontological grounding is fragile or contested. This makes regulatory kinds institutionally enacted simulations of kindhood—classifications that acquire the appearance of stability and epistemic utility through recursive deployment rather than ontological grounding.

They are socially constructed, causally consequential, and epistemically portable—but always underwritten by recursion, not essence.²

Although race will be examined as a regulatory kind in Section 5, this application should not be mistaken for an endorsement of race’s coherence as a biological or epistemic kind.³ The account aims to model institutional behavior, not to validate classificatory categories.

In the next section, this framework is developed in detail. Regulatory kinds are defined by their structural and functional properties, and positioned as a distinct metaphysical category within the broader typology of kinds in science and society. In the following section, they are introduced as a distinct metaphysical category, bridging the gap between social ontology and the philosophy of science by explaining how simulated kindhood emerges through institutional practice.

2. What Is a Regulatory Kind?

This section provides a positive account of regulatory kinds as a distinct metaphysical category. It begins with a formal definition, elaborates the structural features that give these kinds their apparent epistemic power, and distinguishes them from adjacent kind-theoretic models. The aim is to establish regulatory kinds not merely as a subcategory of social

²Not all classifications that appear across institutions qualify as regulatory kinds. The key criterion is recursive epistemic uptake: the category must be treated as explanatory or predictive, not just recorded or archived. For example, the presence of “marital status” in multiple systems does not automatically make it a regulatory kind unless it structures reasoning, intervention, or resource allocation.

³For critiques of the epistemic instability of race classifications—especially in biomedical and public policy domains—see Roberts (2011) and Kahn (2013). These objections are addressed in Section 5.

kinds, but as a necessary metaphysical refinement: a tool for understanding how some social classifications come to *simulate kindhood* across institutional domains.

2.1 Definition and Core Properties

A *regulatory kind* is a socially constructed classification that achieves the appearance of kindhood through recursive institutional uptake. It is stabilized not by intrinsic properties or homeostatic mechanisms, but by repeated use across systems that treat the classification as epistemically salient. These kinds are functionally stabilized and socially sustained: they behave “as if” they were natural kinds because multiple institutions coordinate actions around them.

Regulatory kinds exhibit three core features:

1. **Recursive deployment:** The classification is used across multiple institutional settings—medical, legal, educational, bureaucratic—in ways that implicitly assume the category’s coherence. Each deployment reinforces the classification’s apparent stability and normalizes its inferential use. Recursive uptake functions as a stabilizing mechanism: each application reinforces the category’s perceived reliability and salience, increasing its epistemic portability across contexts.⁴
2. **Cross-contextual portability:** The same category is applied across epistemically distinct domains without local redefinition, as though its referent were stable. “Black” or “Asian” as racial categories, for example, appear in genetic research, census reports, healthcare policy, and crime data, often without alignment or justification.
3. **Apparent epistemic utility:** Regulatory kinds are used not merely to record but to explain, predict, allocate, and intervene. They are part of institutional inference practices—whether or not they meet formal standards of scientific reliability.

These features jointly produce what I call *simulated kindhood*—the institutional enactment of classification as if it tracked a stable underlying kind, even when no such natural structure exists. Simulation here does not mean illusion. It denotes the functional role that a classification plays within institutional reasoning—organizing decisions, guiding inference, and structuring resource allocation—without presupposing natural unity or essential grounding.⁵ It is a real-world mechanism of kindhood stabilization—not a heuristic, metaphor, or misperception.

⁴This recursive structure mirrors the sociotechnical stabilization of classification systems discussed in Bowker and Star (1999).

⁵This use of “simulation” is conceptually adjacent to Ereshefsky’s (2018) pragmatic kinds, but differs in scope. Whereas pragmatic kinds are tied to local epistemic purposes, regulatory kinds simulate portability across multiple institutional domains.

2.2 Distinguishing Regulatory Kinds from Other Social Kinds

Regulatory kinds share important traits with other social kinds but diverge in key respects. To sharpen their metaphysical profile, it is useful to distinguish them along three axes:

- **Ontological dependence:** Regulatory kinds are not grounded in natural or intrinsic properties. Unlike Boydian natural kinds, they are not stabilized by causal homeostasis. They do not track real clusters in the world, but recursively stabilized practices.
- **Epistemic function:** Unlike Haslanger’s ideologically saturated kinds, regulatory kinds are not defined primarily by their role in sustaining normative power structures. Although many do participate in those systems, what defines them is their epistemic uptake—how institutions use them to structure knowledge and action.
- **Interactivity:** Unlike Hacking’s looping kinds, regulatory kinds do not require subject recognition or behavioral feedback. The classification may persist regardless of whether individuals identify with, internalize, or resist it. The kindhood is enacted through systems, not persons.

Although many kinds participate in these systems, what defines regulatory kinds is a specific institutional structure: the category is treated by multiple systems *as if* it supports reliable generalizations, even when its ontological grounding is fragile or contested.

Not all socially constructed classifications become regulatory kinds. Some fail to achieve sufficient cross-contextual uptake or lack perceived epistemic utility. Regulatory kinds emerge only when recursive institutional usage stabilizes the classification’s operational role across domains.⁶

Boyd and the Challenge of Accommodation

Boyd’s theory of natural kinds presents the most direct challenge to the regulatory kind framework. For Boyd (1999), kinds are defined by their ability to support successful inductive inferences across contexts, typically due to underlying causal structures. Social kinds may qualify, he argues, provided they meet the “accommodation condition”: their classification must align with stable causal patterns in the world that make the kind scientifically tractable.

⁶For example, some now-abandoned categories such as “phrenic constitution,” “hysteria,” or “sociopathic mother” circulated briefly but failed to gain recursive institutional traction. Without sustained epistemic uptake, classifications tend to fade or fragment.

Regulatory kinds differ crucially. Their inductive usefulness is often tenuous, fragmentary, or context-bound. Their utility is not evidence of deep natural structure but a product of classification’s recursive reinforcement. As such, they fall short of Boydian natural kind status. Yet their apparent explanatory power remains—embedded in institutional usage. Regulatory kinds occupy the gray zone between heuristic convenience and metaphysical legitimacy: they behave *as if* they track causal structure, even when they do not.

Haslanger and the Function of Ideology

Haslanger (2012) offers a powerful account of kinds shaped by political domination. Race and gender, on her view, are *ideologically saturated*: they exist in part because they help reproduce structural inequalities. While many regulatory kinds overlap with these categories (and often exacerbate structural injustices), the concept of a regulatory kind is not defined by its ideological function. A regulatory kind may persist even in the absence of active norm enforcement, provided it continues to perform epistemic work within institutions.

In this respect, regulatory kinds emphasize structure over ideology. They may be compatible with Haslanger’s view but do not reduce to it. They capture the ways in which classifications acquire a life of their own—circulating across systems, generating data streams, and producing real-world effects even when their ideological content is diluted or obscured.

Hacking and the Limits of Looping

Hacking’s looping kinds are among the most widely cited models of human classification. His work captures the dynamic relationship between labels and behavior: when people internalize or resist classifications, the kind itself evolves. But many classifications that structure institutional life are not subject to looping feedback. Race, for instance, may shape how a patient is treated in an emergency room, regardless of how they identify or whether they accept the label. Similarly, data-driven risk assessments often apply racial or economic proxies mechanically—without regard for individual response.

Regulatory kinds differ from Hacking’s looping kinds in a critical respect: they do not rely on subject recognition or behavioral self-adjustment. While looping kinds presuppose that classification changes behavior in ways that recursively affect the category itself, regulatory kinds operate primarily through normative-expectational feedback: institutions expect the classification to function, and design systems accordingly. The subject may remain unaware, indifferent, or resistant — yet the classification persists because it is infrastructurally necessary. In this sense, regulatory kinds simulate kindhood not through interactive dynamics,

but through the alignment of institutional norms, expectations, and outputs.

Regulatory kinds highlight how classification can be effective *without interactivity*. They are not necessarily responsive to the behavior of those classified. They derive stability not from subjectivity, but from recursive circulation among institutions.

2.3 Simulated Kindhood and Institutional Epistemology

The central contribution of this model is to theorize *simulated kindhood*: the institutional enactment of classification as if it were a natural kind. This simulation has epistemic consequences. Once a classification is treated as predictive or explanatory, it becomes embedded in systems that rely on it—risk modeling, triage protocols, eligibility systems. It creates feedback loops of a different kind: not loops of identity or behavior, but of institutional reinforcement.

Crucially, this model is compatible with conceptual instability. The referent of the classification may shift over time or differ across contexts. But as long as the category continues to structure institutional action, it maintains the appearance of kindhood. This helps explain why race, disability, and criminality often appear in scientific and bureaucratic discourse with the same rhetorical and epistemic weight as more stable categories, even when their ontological basis is contested.

This framework is not intended to justify or rehabilitate contested classifications such as race. Rather, it is designed to explain how such categories persist and function across systems even when their coherence is widely rejected.⁷

To further distinguish regulatory kinds from nearby metaphysical models and clarify their functional criteria, I now offer a quasi-formal characterisation.

2.4 A Quasi-Formal Characterisation of Regulatory Kinds

The discussion so far has introduced regulatory kinds as an institutional ontology — a category of classification that simulates kindhood through recursive uptake and infrastructural embedding. To make this account tractable and distinguishable from existing metaphysical models, we can now offer a quasi-formal schema.

Let C be a classification, and let $I = \{i_1, i_2, \dots, i_n\}$ be a non-trivial set of institutional domains (e.g., medicine, forensics, census, education). Then C qualifies as a *regulatory kind* just in case the following three conditions hold:

⁷For critiques of the epistemic instability of race classifications—especially in biomedical and policy domains—see Roberts (2011) and Kahn (2013). These objections are taken up directly in Section 5.

1. **Recursive Institutional Uptake:** C is adopted, transmitted, or instantiated across multiple $i \in I$, such that the classificatory output of i_a is recursively reinterpreted or operationalized in i_b .
2. **Normative-Expectational Feedback:** Each $i \in I$ treats C as sufficiently stable to underwrite inference, intervention, or coordination. This feedback may be statistical, normative, administrative, or computational.
3. **Epistemic Entrenchment:** C becomes encoded into institutional tools, infrastructures, and routines — databases, diagnostic forms, eligibility criteria — such that its future use is path-dependent, even if its empirical basis is unstable.

This structure allows regulatory kinds to simulate the epistemic roles of natural kinds without appealing to causal unity, semantic precision, or subjective recognition. Unlike Hacking’s interactive kinds, regulatory kinds do not require self-identification or looping effects. Unlike Haslanger’s ideologically saturated kinds, they need not function as vehicles of critique. Their persistence emerges not from endorsement or resistance, but from recursive instrumentalization across institutional systems.

This quasi-formal account also clarifies the role of *simulated kindhood* in practice: a category C simulates kindhood when it behaves, institutionally, like a kind — not because it tracks natural divisions, but because systems expect it to.

2.5 Summary and Justification

Regulatory kinds are not a new name for social kinds in general. They are a specific metaphysical model for explaining how certain socially constructed classifications acquire simulated kindhood across institutions. They do not depend on natural structure, ideological saturation, or looping interactivity. They depend on recursive epistemic uptake: systems use the classification to organize, infer, allocate, and decide—thereby giving it durability, portability, and apparent coherence.

But why introduce regulatory kinds as a distinct metaphysical category, rather than treating them as subtypes of existing kind models? Why not regard them as ideologically saturated kinds with epistemic features, or as Boydian kinds with attenuated accommodation?

The answer is explanatory scope. None of the established frameworks fully accounts for classifications that:

- lack subject recognition or behavioral feedback (contra Hacking),
- operate independently of normative ideology (contra Haslanger),

- and do not satisfy causal accommodation or ontological depth (contra Boyd or Khalidi),

yet still circulate across domains, structure institutional behavior, and guide reasoning and resource distribution.

Regulatory kinds are introduced to capture this unique pattern of kind-like behavior: the recursive simulation of kindhood by institutions that depend on epistemic regularity, even in the absence of natural or ideologically coherent grounding. Existing theories explain identity formation, political legitimation, and scientific inference. Regulatory kinds explain how classifications that are fragmented, contested, or empirically unstable nevertheless function in policy, law, medicine, and administration as if they were reliable kinds.

Without this category, we risk obscuring the mechanism by which unstable classifications become materially entrenched and epistemically routinized. Regulatory kinds offer a metaphysical tool for understanding why these classifications endure—not by virtue of what they are, but by virtue of how they are used.

Regulatory kinds, in this sense, are constructivist realist kinds: they are not nominal groupings or rhetorical artifacts, but institutionally enacted structures that produce real epistemic and material consequences.⁸

The next section develops this account of simulated kindhood in greater detail, with attention to its epistemic structure and empirical consequences.

3. Simulated Kindhood

The distinctive contribution of regulatory kinds is not that they describe socially constructed classifications, but that they explain how such classifications come to function epistemically like kinds, even in the absence of causal unity or natural kindhood. This section develops the concept of *simulated kindhood*: the recursive, institutionally enacted performance of kind-like behavior by unstable or ontologically shallow categories. It is this phenomenon that regulatory kinds are introduced to model.

3.1 Simulation Without Illusion

To say that a classification simulates kindhood is not to suggest error, deception, or confusion. Simulated kindhood is not an illusion; it is a structured form of epistemic enactment. Institutions treat a classification as if it were a kind—using it to generalize, predict, allocate, and coordinate—and through this use, the classification accrues the functional profile of a

⁸Regulatory kinds are compatible with pluralist accounts of social kindhood, such as Spencer’s (2019) or Khalidi’s (2016). They do not deny the legitimacy of more localized or referential kinds. Rather, they highlight a broader structural pattern: simulation of kindhood across epistemic contexts through recursive institutional enactment.

kind. The simulation is real in its effects, even if the kindhood is constructed rather than discovered.

This sharply distinguishes simulated kindhood from mere heuristic usage. Heuristics are acknowledged approximations: short-cuts in reasoning designed to cope with complexity or uncertainty. Simulated kinds, by contrast, are often treated as stable, tractable classifications. Their apparent stability is not marked as provisional. It is assumed. Institutions build around the classification as though it were explanatorily robust. Diagnostic criteria, eligibility rules, legal definitions, and data schemas all reify the simulation by embedding it into practice.

9

3.2 The Epistemic Mechanics of Simulation

Simulated kindhood involves a distinctive set of epistemic functions. These include:

- **Generalization:** treating members of the classification as predictable bearers of shared traits (e.g., “Black patients tend to respond differently to medication X”).
- **Coordination:** using the classification to sort individuals across systems (e.g., education, public health, criminal justice).
- **Prediction:** deploying the classification in algorithmic or statistical models (e.g., recidivism risk tools, actuarial assessments).
- **Explanation:** invoking the classification to make causal sense of disparities (e.g., racial gaps in outcomes, health inequities).

These functions do not require natural unity. They require stability in use. Simulation occurs when recursive institutional uptake makes a classification epistemically legible and materially consequential. The more a system depends on a category for its inferential routines, the more that system reinforces the apparent kindhood of the category.

3.3 Case Illustration I: Race and Diagnostic Modeling

Consider the use of race in nephrology. For decades, estimated glomerular filtration rate (eGFR) equations in American medicine included a race modifier: the value for “Black” patients was adjusted upward, based on assumptions about average muscle mass. This

⁹Unlike nominalism, which treats classifications as arbitrary labels, simulation explains how classifications acquire kind-like features through epistemic embedding. Unlike heuristic categories, which are explicitly provisional, simulated kinds are treated as epistemically real—structuring institutional inference as though grounded in natural unity.

practice was justified on the grounds of statistical generalization. But the classification “Black” was not biologically coherent, nor did it reflect any consistent causal mechanism. It simulated kindhood: race appeared to carry diagnostic meaning because institutional practice treated it as if it did.

This simulation was recursive. Medical guidelines encoded the adjustment. Electronic health records maintained race fields. Physicians came to expect racial modifiers. The more the classification was used, the more its apparent medical relevance was reinforced. And yet this simulation produced harm—delaying treatment, misrepresenting individual variation, and obscuring socioeconomic causality. The classification endured, not because it was explanatory, but because it was embedded. Its kindhood was simulated, not discovered.

3.3.1 Case Illustration II: BiDil and Regulatory Reification

Another powerful example is BiDil, the first FDA-approved race-specific medication, indicated for “self-identified Black patients with heart failure.” As Kahn (2013) documents, this approval was driven not by genetic evidence, but by regulatory convenience: race functioned as a proxy for unmeasured social variables. The drug’s efficacy data could not justify racial specificity on biological grounds. But race simulated kindhood by enabling coordination—among drug trials, insurance approvals, and pharmaceutical marketing. Once adopted, the race tag became part of the drug’s identity. This is simulation in its clearest form: a classification treated as a causal kind, sustained not by biology but by policy, precedent, and institutional inertia.

3.4 Simulation as Institutional Epistemology

Simulated kindhood is best understood as a feature of institutional epistemology: the way systems generate, stabilize, and act on categories in order to manage populations and uncertainty. In environments of high complexity—public health, social services, law enforcement—systems require classifications that support coordination and inference. Kindhood is simulated when the need for systemic regularity outweighs the metaphysical legitimacy of the classification.

What results is a feedback structure:

classification → institutional uptake → epistemic embedding → policy coordination → recursive reification

This is not a feedback loop of self-identification (as in Hacking), but a loop of data, expectation, and infrastructure. The classification appears stable because institutions act

on it as if it were. This explains how categories that are epistemically fragile or ethically contested nevertheless persist and function across domains.

The kind-like behavior of race becomes particularly visible when tracing its uptake across distinct but interacting institutional domains. In forensic anthropology, practitioners routinely estimate racial ancestry from skeletal remains to aid identification — a practice that persists despite philosophical and biological objections to the coherence of race. In clinical medicine, race is embedded into decision algorithms, such as the use of race-adjusted equations for kidney function (eGFR), which alter diagnostic thresholds for Black patients. Meanwhile, public policy mandates race reporting through fixed federal categories (e.g., OMB Directive 15), which standardize classification across education, housing, and health institutions. What links these domains is not semantic continuity or referential precision, but institutional feedback: once a classification like “Black” is encoded across systems, it begins to behave as if it were a stable kind — explanatory in forensics, predictive in medicine, and administratively actionable in policy. This cross-domain uptake is what allows race to simulate kindhood and sustain epistemic traction despite ontological instability.

3.5 Simulation and Critical Race Epistemology

Simulated kindhood is compatible with, and in many ways explains, the critiques advanced by critical race theorists such as Dorothy Roberts and Jonathan Kahn. These authors emphasize that the institutional salience of race is often a product of flawed data practices, legacy classification schemes, or political convenience—not biological validity. Roberts (2011), for instance, critiques epigenetics studies that link race to DNA methylation as evidence of embodied inequality, not natural clustering. Kahn (2013) shows how race is made to seem causally meaningful through institutional procedures like FDA labeling or NIH funding mandates. These analyses do not undermine the concept of simulated kindhood. They help illustrate it. The present framework is not an endorsement of such classifications, but an account of how they acquire functional stability through epistemic circulation.

10

3.6 Failure, Fragility, and Simulation Limits

Simulated kindhood is not always successful. Classifications can fail to stabilize, produce epistemic distortions, or face breakdown under contestation. For example, when categories are too inconsistent across domains, simulation falters. If “Latinx,” “Hispanic,” and “Spanish-speaking” are variably defined in education, health, and immigration policy, recur-

¹⁰As Kahn writes, “The categories persist, not because they are true, but because they are administratively convenient and epistemically entrenched.” (Kahn 2013, p. 88).

sive uptake becomes disjunctive rather than reinforcing. Simulation requires consistency of use, even if the underlying referent is unstable.

Likewise, simulated kindhood can backfire. It may obscure real causes, justify inequities, or reinforce spurious inferences. When a classification is overfitted to institutional needs, it becomes functionally indispensable but epistemically toxic. Understanding this risk is part of what regulatory kinds help explain. They are not normative endorsements. They are metaphysical models of how categories acquire kind-like behavior in practice, often despite their instability or injustice.

3.7 Summary: The Lifecycle of Simulated Kindhood

Simulated kindhood is the epistemic mechanism through which regulatory kinds operate. It explains how institutions come to treat socially constructed, ontologically shallow categories as if they were natural kinds. The simulation is not merely cognitive or rhetorical. It is infrastructural: embedded in records, forms, protocols, and procedures. Through recursive uptake and epistemic reinforcement, simulation makes a classification appear stable, generalizable, and predictive.

This framework provides a new account of kindhood in applied institutional contexts. Rather than beginning with ontological commitments, it begins with epistemic behavior: how systems come to depend on and reproduce certain categories. Regulatory kinds simulate kindhood not by tracking nature, but by structuring inference.

In the next section, this framework is applied to the classification of race. The goal is not to defend race as epistemically coherent, but to analyze how it functions across medical, legal, and policy domains as a regulatory kind.

4. Race as a Regulatory Kind

The category of race constitutes a critical test case for the framework of regulatory kinds. If the concept is to have explanatory power, it must be capable of accounting for classifications that are widely recognized as constructed, inconsistent, and politically charged, yet remain central to institutional decision-making and epistemic practice. Race is exemplary in this regard. It is used across medical, legal, educational, demographic, and forensic domains to sort, assess, allocate, and intervene—even as its ontological legitimacy is repeatedly challenged.

This section examines how racial classification functions across institutional settings, not to argue for its coherence or legitimacy, but to show how it behaves like a regulatory kind: a socially constructed classification that simulates kindhood through recursive uptake and epistemic embedding. Drawing on empirical examples, the analysis demonstrates how

race achieves apparent cross-contextual stability, how it is used to generate generalizations and guide inferences, and how its simulated kindhood persists despite deep conceptual and normative instability.

4.1 The Puzzle of Racial Persistence

Race poses a conceptual challenge that lies at the heart of regulatory kind theory. It is among the most widely used social classifications in modern institutions—yet also among the most widely contested. Most philosophers of race agree that racial categories lack biological coherence, stable semantics, or essentialist metaphysical foundations.¹¹ Nevertheless, racial classifications are deeply embedded in the epistemic infrastructure of contemporary life: they inform medical diagnostics, structure public health responses, organize demographic data, guide risk modeling, and regulate legal and financial systems.

This presents a puzzle. How can race be simultaneously unstable and indispensable? How can a category so thoroughly deconstructed by critical scholarship remain epistemically active—treated as if it supported generalization, prediction, and explanation? Why do institutions continue to act on racial classifications as though they were tracking meaningful divisions in the world?

This section addresses that puzzle by showing how race functions as a *regulatory kind*: a classification that simulates kindhood through recursive institutional uptake. The argument is not that race is real in a biological or metaphysical sense. Nor is it that race is epistemically justified in every context. Rather, it is that race has become functionally stabilized and epistemically routinized across institutions. This simulated kindhood explains race’s institutional persistence—not because it tracks causal structure, but because systems treat it *as if* it does.

The analysis that follows is deeply empirical. It draws on examples from medicine, public health, law, demography, and genomics to show how race is enacted across domains. In doing so, it responds directly to critiques that emphasize the incoherence or normative danger of racial classification. These critiques are important—and often justified. But they do not explain why racial categories persist. Regulatory kind theory aims to fill that gap.

4.2 Recursive Uptake Across Institutional Domains

To function as a regulatory kind, a classification must be used across multiple institutional domains, with sufficient consistency to simulate kindhood. In this section, I survey how racial categories are deployed in medicine, public health, genomics, forensic science, law, and demography. The aim is not to argue that these uses are epistemically justified, but

¹¹See Glasgow (2009); Haslanger (2012); Spencer (2019); Roberts (2011).

to show how they achieve a form of functional stability through recursive uptake: each institutional application reinforces the appearance of coherence and utility, producing the illusion of a kind.

Medicine and Public Health. Racial classification is deeply embedded in clinical practice. One of the most widely cited examples is the use of a race modifier in estimated glomerular filtration rate (eGFR) equations. Until recently, “Black” patients in the U.S. were assigned higher baseline kidney function scores than non-Black patients, based on assumptions about average muscle mass. This adjustment was encoded in medical guidelines, electronic health record systems, and laboratory algorithms. The classification appeared to carry diagnostic weight—not because of biological evidence, but because institutions acted on it as if it did.¹²

Similarly, spirometry devices often apply race-based correction factors, lowering expected lung function for “Black” and “Asian” patients. These modifiers are legacy artifacts of 19th-century race science, yet persist in modern devices and training protocols.¹³ Here too, the race variable simulates kindhood by enabling inference: physicians act on it, adjust treatment thresholds, and enter it into databases. Whether or not the category tracks any stable biology is beside the institutional point.

Genomics and Biomedical Research. In clinical trials and population studies, racial categories are used as groupings for stratification and analysis. NIH funding guidelines explicitly require researchers to disaggregate data by race and ethnicity.¹⁴ This institutional requirement ensures that race becomes epistemically active: even if researchers recognize that race is a poor proxy for genetic ancestry, it is nonetheless used to organize findings, allocate funding, and signal compliance.

Genomic research also enacts race through continental population clusters (e.g., “African,” “East Asian,” “European”), which are treated as proxies for ancestry. These groupings often reimport U.S. census categories into global research contexts, despite radically different local meanings.¹⁵ Again, race simulates kindhood—not because it reflects deep population structure, but because it serves as a unit of organization and generalization.

Forensic Science and Risk Modeling. Racial classifications also operate in criminal justice algorithms and forensic assessments. Tools like COMPAS, used to predict recidivism

¹²Braun, L. (2014). *Breathing Race into the Machine*. University of Minnesota Press. See also Roberts (2011), ch. 5.

¹³Braun et al. (2007). “Racial Categories in Medical Practice.” *PLOS Medicine*, 4(9).

¹⁴See: NIH Revitalization Act, 1993.

¹⁵Fujimura, J.H. & Rajagopalan, R. (2011). “Different Differences.” *Social Studies of Science*, 41(1).

risk, explicitly or implicitly incorporate race as a variable, often through correlated proxies like zip code, employment history, or educational attainment.¹⁶ Racial categories also appear in forensic anthropology, where skeletal remains are assigned racial labels based on cranial features. These assignments are fed into criminal databases and used in identification procedures.¹⁷

In both cases, race operates as if it were a reliable, inductively powerful kind. Whether these applications are empirically justified or normatively acceptable is not the point. What matters is that the classification is epistemically routinized: it enters algorithms, policy decisions, and evidentiary structures with the weight of a kind.

Demography, Education, and Policy. Census categories codify race for the purposes of population tracking, legislative redistricting, and federal funding. In the United States, OMB Directive 15 establishes five official racial categories that are used by nearly all federal agencies, including the Department of Education, Housing and Urban Development, and the Department of Justice.¹⁸ These categories appear in school records, employment statistics, health disparity reports, and more.

Despite lacking scientific basis, these categories structure access to services, signal legal compliance, and serve as the basis for social scientific research. The portability of racial categories across these domains—without local redefinition—further simulates kindhood. The same label (“Black,” “Asian,” “White”) is used across vastly different systems to infer risk, track outcomes, and guide intervention.

Epigenetics and the Embodiment of Inequality. Recent work in social epigenetics suggests that structural racism produces biological effects, such as altered gene expression through DNA methylation pathways.¹⁹ These studies show how racism “gets under the skin”—but often do so by reifying race as a biological exposure category. While the intention is to highlight inequality, the result is often a recursive uptake of race in biomedical discourse. As Roberts (2011) warns, this risks re-biologizing race by simulating kindhood at the molecular level, even when the underlying cause is social.

¹⁶Angwin et al. (2016). “Machine Bias.” *ProPublica*.

¹⁷Sauer, N. (1992). “Forensic Anthropology and the Concept of Race.” *American Anthropologist*, 94(1).

¹⁸Office of Management and Budget (1997). “Revisions to the Standards for the Classification of Federal Data on Race and Ethnicity.”

¹⁹Kuzawa, C. & Sweet, E. (2009). “Epigenetics and the Embodiment of Race.” *American Journal of Human Biology*, 21(1).

4.3 Apparent Epistemic Utility

Institutional systems do not merely record racial categories; they use them to organize knowledge, structure inference, and justify intervention. This functional role gives racial classification the appearance of kindhood. Although race may not track natural properties or intrinsic causal structures, it is nonetheless treated as a basis for explanation, prediction, and allocation. This apparent utility is not incidental — it is constitutive of regulatory kind behavior.

Explanation and Risk Framing. Racial categories are routinely invoked to explain outcome disparities in medicine, education, and social policy. Epidemiological reports describe disease prevalence by race. Public health studies identify “risk factors” associated with racial groups. For example, Black women in the U.S. are said to be three to four times more likely to die from pregnancy-related complications than White women — a statistic repeated in policy documents, media coverage, and institutional reform proposals.²⁰

While such statistics are often motivated by anti-racist aims, they also function epistemically: race becomes a variable that structures causal narratives. The category is treated as if it were a meaningful explanatory unit — not merely a descriptor of exposure, but a site of attribution. Even when researchers qualify these associations as social or structural, institutional logic tends to treat the racial label itself as actionable knowledge.

Prediction and Algorithmic Sorting. Race is increasingly used in predictive systems — from actuarial models to machine learning tools. In healthcare, algorithms may incorporate racial data to estimate risk for conditions like diabetes, stroke, or COVID-19 severity.²¹

In criminal justice, tools like COMPAS predict recidivism risk using variables highly correlated with race. These models are used to inform bail, sentencing, and parole decisions — effectively treating race-adjacent data as a proxy for future behavior.²²

In both domains, race is not merely descriptive. It is an active input in epistemic infrastructure: used to produce forecasts, trigger interventions, and automate decision-making. The category behaves like a kind — not because it captures stable causal relationships, but because it is treated as if it does.

Allocation and Eligibility. Racial categories also guide resource distribution. In public health, targeted programs use race to prioritize vaccine outreach, disease screening, and

²⁰CDC (2022). “Racial and Ethnic Disparities Continue in Pregnancy-Related Deaths.” Centers for Disease Control and Prevention.

²¹Obermeyer, Z. et al. (2019). “Dissecting racial bias in an algorithm used to manage the health of populations.” *Science*, 366(6464).

²²Angwin, J. et al. (2016). “Machine Bias.” *ProPublica*.

intervention funding. For example, the U.S. Department of Health and Human Services has allocated COVID-19 response funds partly based on racial impact data — using group-level classifications to identify “disproportionately affected communities.”²³

In education, racial demographics affect school funding formulas, diversity incentives, and admissions review processes. Institutions rely on aggregate racial data to justify reallocation, affirmative action, and equity measures. These practices are grounded in moral and political aims — but they also functionally reinforce the category’s kind-like status. Race becomes a criterion that structures who gets what, when, and how.

Epistemic Consequences. Across these domains, the institutional use of race generates the appearance of coherence and predictive reliability. This is the core of simulated kindhood: a category that acquires the surface-level traits of a kind — generalizability, inferential productivity, policy salience — without satisfying the ontological conditions of kindhood. It appears to work because institutions have built around it.

This appearance is recursive. The more race is used to explain, predict, or allocate, the more empirical studies reference race as a meaningful variable. Policy documents cite those studies. Data platforms encode the categories. Algorithms are trained on them. And the cycle continues — often without confronting the instability, heterogeneity, or arbitrariness of the classification itself.

To be clear, this is not a claim about the correctness of racial data use. It is a claim about institutional epistemology: systems act as if race were a reliable, explanatory, and generalizable kind — and in doing so, reinforce its epistemic reality.

4.4 Semantic Instability and Regulatory Resilience

A common objection to treating race as a kind — even in simulated or regulatory terms — is that racial categories are conceptually incoherent and empirically inconsistent. The category “Black,” for instance, encompasses vastly different populations across the U.S., Brazil, South Africa, and the U.K., each shaped by unique historical, cultural, and political dynamics. The referents of racial labels shift across time and space. Even within a single national context, definitions are contested, overlapping, and subject to strategic deployment.²⁴

Critics such as Roberts (2011) and Kahn (2013) have rightly emphasized this instability. They argue that racial categories are too fragmented to serve as reliable units of scientific analysis — and that apparent utility is often a result of flawed methodologies, institutional inertia, or socio-political convenience. These critiques are not only valid, but essential.

²³HHS Office of Minority Health (2021). “Advancing Health Equity through the American Rescue Plan.”

²⁴Morning, A. (2011). *The Nature of Race*. University of California Press.

Yet they do not fully explain how and why race continues to be treated as epistemically meaningful across institutional contexts.

Regulatory kind theory offers a complementary perspective. It acknowledges semantic instability as a feature, not a bug — and explains how classifications like race can persist despite it. Regulatory kinds do not require internal coherence or causal unity. They require only that institutions treat the category as if it were a kind: generalizable, portable, and functionally useful.

The key mechanism here is recursive uptake. When disparate institutions — health systems, census bureaus, courts, and research bodies — coordinate their practices around the same category label, the classification begins to simulate epistemic coherence. The label “Black” may refer to different things in different settings, but its recursive use creates the appearance of stability. This enables the category to travel across domains with minimal reinterpretation, producing what sociologists call “institutional isomorphism.”²⁵

In this sense, the resilience of racial classification is not a sign of its semantic clarity, but of its epistemic embedding. Systems need consistent units to manage populations, allocate resources, and justify interventions. Race functions as such a unit not because it is conceptually well-formed, but because it is infrastructurally entrenched.

This distinction is crucial. Regulatory kinds are not metaphysically stable; they are institutionally stabilized. They persist not by tracking a single referent, but by anchoring action. As long as the classification produces usable outputs — in statistics, policies, or procedures — it retains its simulated kindhood. Instability at the level of meaning does not undermine its operational role. On the contrary, the abstraction and flexibility of racial categories may enhance their institutional utility.

The BiDil case offers a striking example. Approved for “self-identified Black patients,” the drug’s labeling sidestepped the question of what “Black” means. No genetic or biological criteria were provided. The label’s effectiveness depended not on semantic precision, but on institutional compatibility: FDA approvals, marketing strategies, and insurance reimbursement protocols all accepted the racial tag as functionally sufficient.²⁶ The category’s resilience was not epistemic in the traditional sense. It was regulatory.

This resilience explains why race persists in genomic studies, clinical guidelines, public health campaigns, and educational metrics — even as its meaning remains contested. The classification continues to function as a regulatory kind because institutions act on it as if it were epistemically robust, regardless of its ontological status.

²⁵DiMaggio, P. & Powell, W. (1983). “The Iron Cage Revisited.” *American Sociological Review*, 48(2).

²⁶Kahn, J. (2013). *Race in a Bottle*. Columbia University Press.

4.5 Race Compared to Other Kind Models

The regulatory kind framework does not exist in a vacuum. It builds on — and diverges from — several influential accounts of social classification. In this section, I briefly compare the regulatory kind model to three prominent alternatives: Hacking’s interactive kinds, Haslanger’s ideologically saturated kinds, and Boyd’s natural kinds via homeostatic property clusters. The aim is not to dismiss these models, but to clarify what the regulatory kind framework adds to our understanding of race’s institutional persistence.

Hacking and Interactive Kinds. Hacking (1995) introduced the concept of “looping kinds” to explain classifications that interact with the people they label. When individuals internalize, resist, or modify their behavior in response to being classified — as with psychiatric diagnoses or categories of deviance — the classification itself evolves. This interactivity creates a dynamic epistemic loop.

Race certainly exhibits looping effects in many contexts. However, this alone cannot explain its persistence in systems that do not depend on self-identification or behavioral feedback. In medical diagnostics, forensic databases, or algorithmic models, race is often assigned externally and functions independently of personal uptake. Regulatory kinds fill this explanatory gap. They show how race can be recursively enacted across institutions, even in contexts where subject response is irrelevant.

Haslanger and Ideologically Saturated Kinds. Haslanger (2012) offers a compelling account of social kinds as structured by ideology and power. On her view, classifications like race and gender are vehicles for oppression — they persist because they serve dominant social interests and reinforce unjust hierarchies.

This perspective is crucial for understanding the normative dimensions of racial classification. But it does not fully explain why race functions epistemically as a kind in contexts that are not overtly ideological — such as bioinformatics pipelines, diagnostic software, or census tabulation. Regulatory kinds do not displace Haslanger’s account; they complement it by capturing the ways in which race becomes epistemically routinized in systems that may be politically neutral in appearance, yet materially consequential in effect.

Boyd and Natural Kinds. Boyd (1991) famously argued that natural kinds support inductive reliability because their members share a causal structure that satisfies the accommodation condition. A kind is natural, in this sense, if it underwrites successful scientific generalizations and explanations due to underlying property clusters.

Race, on this view, is not a natural kind — and few philosophers of science would claim otherwise. Yet race continues to be used to guide generalization and explanation in

institutional settings. Regulatory kinds explain this puzzle. They simulate kindhood by accumulating epistemic infrastructure: forms, algorithms, risk schemas, eligibility codes. These constructs do not reflect natural structure, but create inferential stability through practice.

Spencer and Referentialist Pluralism. Spencer (2014, 2019) defends a referentialist approach to race, arguing that terms like “race” refer to distinct groupings in different contexts — e.g., U.S. census categories versus biomedical population clusters. This pluralism accommodates local variation without requiring a single theory of race.

Regulatory kinds are compatible with this pluralism but offer a different emphasis. They are not theories of racial semantics. They are metaphysical models of how categories function institutionally across contexts, even when their referents shift. Regulatory kinds do not claim that “Black” means the same thing everywhere. They claim that “Black” is used as if it does — and that this usage stabilizes institutional inference.

Conclusion. What distinguishes regulatory kinds is not their ideological origin, semantic content, or interactive character — but their institutional function. They explain how a contested and ontologically unstable category like race can persist, appear epistemically tractable, and structure social action across multiple domains. Regulatory kinds simulate kindhood not by reflecting the world, but by organizing it.

4.6 Institutional Naturalization

If the preceding analyses converge on a single insight, it is this: race has become naturalized not through biological discovery, but through institutional function. Its apparent stability, portability, and predictive power arise not from intrinsic properties, but from recursive usage across systems that rely on it to structure inference, coordinate action, and manage populations.

This functional naturalization is what regulatory kind theory helps to explain. The notion of simulated kindhood captures how race, despite its ontological fragility and semantic inconsistency, accrues the surface-level epistemic features of a kind — generalizability, explanatory traction, administrative utility. These features are not metaphysically inherited; they are socially and materially enacted.

Each institutional domain examined — from medicine and genomics to criminal justice and public policy — reuses race in ways that reinforce its perceived coherence. These uses do not align neatly; they overlap, diverge, and occasionally contradict one another. Yet the classification persists. Its durability lies not in what it is, but in what it enables. This is the hallmark of a regulatory kind.

Crucially, none of this analysis entails endorsement. To say that race behaves like a kind is not to suggest that it should. It is to note that the category performs epistemic work — sometimes efficiently, often problematically — because it has been built into the architecture of institutional reasoning.

As we move forward, the focus shifts from race in particular to the broader implications of regulatory kinds. What does simulated kindhood reveal about how institutions manage uncertainty? What risks emerge when classificatory structures become epistemically entrenched? And how might regulatory kinds illuminate both the power and the limits of classification as a mode of governance?

5. The Reach and Risks of Regulatory Kinds

The preceding sections have developed and applied the concept of regulatory kinds, showing how a contested classification like race can come to function epistemically across institutional domains. But race is not the only category that behaves this way, nor is epistemic function always benign. Regulatory kinds do not simply describe a phenomenon — they raise questions about how institutions generate, stabilize, and act on classifications whose kindhood is simulated rather than discovered.

This section turns from analysis to reflection. It explores the broader implications of regulatory kinds for institutional epistemology, policy design, and the ethics of classification. It asks: What happens when simulated kindhood is mistaken for natural structure? What risks arise when classificatory tools ossify into objects of belief or decision rules? And how might we recognize regulatory kinds at work beyond race — in domains such as criminality, disability, or economic identity?

The goal is not to undermine the explanatory value of regulatory kinds, but to understand their institutional logic — both what they enable and what they obscure. If regulatory kinds help us make sense of classificatory persistence under metaphysical fragility, they may also help us see how categories come to govern the world, long after we’ve stopped believing in them.

5.1 Institutions and Epistemic Dependence

Institutions require classificatory order to function. In domains such as health, education, justice, and social policy, vast and heterogeneous populations must be rendered legible through the imposition of categories. These categories serve as epistemic shortcuts — tools for prediction, coordination, and resource distribution. They allow systems to act under conditions of partial information and structural uncertainty. But the classifications most

available for this work are often socially constructed, ontologically unstable, and epistemically thin. The institutional solution, frequently, is not to refine the classification, but to routinize it. This is the terrain on which regulatory kinds emerge.

Consider again race in clinical practice. The continued use of race in estimating kidney function (eGFR), despite persistent challenges to its biological validity, was less an epistemic endorsement than a bureaucratic default. Laboratories, insurers, and electronic health record systems had built processes around the race-modified formula. Abandoning the category would have disrupted care pathways, altered patient eligibility, and invalidated historical benchmarks.²⁷ Race was retained not because it was true, but because it was infrastructurally entrenched.

Similar dependencies are visible in education, where race is used to track performance gaps, evaluate funding equity, and implement diversity policies. In such contexts, the epistemic function of race is not to describe a metaphysical reality, but to enable action across complex systems. Once a classification becomes embedded in institutional reasoning — through compliance protocols, regulatory forms, or data schemas — it accrues operational inertia. It begins to behave like a kind because the system depends on treating it as one.

This institutional dependence explains the resilience of certain social classifications. Institutions cannot afford to rebuild their epistemic infrastructure with every theoretical critique. They require regularity. Regulatory kinds deliver that regularity by simulating stability — not through correspondence with nature, but through recursive uptake. When a category supports relatively consistent outputs across multiple domains, it acquires the appearance of generalizability. It becomes trusted — not because it is true, but because it is portable.

Crucially, this is not a sign of institutional irrationality. In high-stakes, resource-limited settings, the imperative to act often exceeds the capacity to interrogate foundational assumptions. Hospitals need triage criteria. Courts need risk assessments. Public agencies need policy benchmarks. In this context, the simulation of kindhood is a feature, not a flaw. It enables decision-making under epistemic constraint.

Yet this very dependence creates blind spots. Once a classification is stabilized through institutional usage, its epistemic profile becomes difficult to challenge. Feedback loops harden into dogma. Critiques are deflected by operational necessity. As Bowker and Star (1999) note, classifications, once entrenched, often appear “natural, necessary, and inevitable.” Regulatory kinds show how that illusion is produced — and why it is so epistemically durable.

The concept of regulatory kinds thus illuminates a paradox at the heart of institutional epistemology: the categories most necessary for governance are often the least ontologically secure. Yet rather than collapsing under their fragility, such categories persist by becoming indispensable. The appearance of kindhood is manufactured not by evidence, but by

²⁷Roberts, D. (2011). *Fatal Invention*, ch. 5.

dependence.

5.2 Simulation and the Creep Toward Reification

Regulatory kinds originate in institutional necessity. They are constructed classifications that simulate kindhood because systems require them to. Yet over time, this simulation can harden. What begins as a pragmatic structure to manage complexity can become an ontological commitment. Classifications designed to enable action are reinterpreted as reflections of reality. The simulated becomes the discovered. This is the creep toward reification.

The mechanism is subtle. Recursive uptake gives a category epistemic durability. Used in diagnostics, eligibility, and risk modeling, the classification appears reliable. Its apparent stability invites inference. Its portability across domains suggests generality. And as more systems incorporate it into their infrastructure, its origin fades from view. Institutions cease to treat the category as a tool and begin to treat it as a truth.

Race offers a paradigmatic case. As Roberts (2011) and Kahn (2013) document, race-based modifiers in clinical algorithms were often introduced as rough proxies — placeholders for complex, unmeasured variables like access to care or structural inequity. But once embedded in clinical workflows, they acquired an aura of biological realism. Race began to appear as a determinant, not just a correlate. Physicians cited race as a risk factor. Patients were assigned racial labels during intake. And the category reentered research protocols as an explanatory variable. What was originally a functional simulation became an epistemic commitment.

This process is not confined to race. In forensic science, algorithmic policing, or welfare assessment, categories such as “criminal risk,” “dependency,” or “eligibility class” are operationalized, iterated, and eventually reified. The model becomes the measure. Predictive tools trained on past classifications reproduce them. Feedback loops intensify. And institutional actors begin to treat these categories as if they marked real divisions in the world, rather than inherited artifacts of prior decisions.

Reification carries both epistemic and ethical risks. It distorts inquiry by masking the constructed nature of the category. It closes space for contestation, since the classification now appears natural. And it legitimizes interventions — often coercive or discriminatory — on the basis of what are, in effect, simulations. Regulatory kinds enable institutions to function under constraint, but they also risk entrenching the very inequalities they purport to manage.

Importantly, this is not a call for epistemic purism. In many cases, institutions cannot wait for metaphysical certainty. They must act. But recognizing the provisional, constructed, and instrumental nature of regulatory kinds is essential to resisting reification. It is what allows systems to remain corrigible — open to revision, contestation, and reflection.

Simulated kindhood is not a flaw. It is a response to constraint. But when simulation is mistaken for truth, regulatory kinds become dangerous. They cease to serve the institutions that created them, and begin to govern them in return.

5.3 Beyond Race: Toward a Taxonomy of Regulatory Kinds

The case of race demonstrates how a socially constructed and epistemically unstable classification can come to function as a kind through institutional recursion. But race is not the only classification to behave this way. The logic of regulatory kinds is general. It applies wherever institutions operationalize uncertain or contested categories to stabilize action, enable inference, and coordinate across domains. Identifying these cases helps clarify the scope — and the limits — of the framework.

One candidate is *criminality*. Risk assessment tools used in pretrial detention, parole decisions, and sentencing — such as COMPAS or PSA — classify individuals into risk tiers based on composite variables. These risk labels simulate kindhood: they guide intervention, justify allocation of surveillance or leniency, and produce recursive data feedback. But they do not track a stable behavioral essence. Rather, they operationalize a probabilistic construct that gains institutional solidity through repeated use.²⁸

Disability offers another example. Legal and medical systems define disability through different logics — functional impairment, diagnostic criteria, or impact on daily life. Yet despite this pluralism, institutional actors must act on the classification: it determines access to accommodations, public benefits, and workplace protections. The category persists not because it maps cleanly onto a biological or psychological kind, but because systems depend on it to regulate eligibility, entitlements, and rights.²⁹

Economic classifications such as *working poor*, *middle class*, or *essential worker* similarly function as regulatory kinds. They are invoked to justify tax policy, wage subsidies, housing programs, or pandemic aid. These labels are fuzzy, contested, and deeply historical — but they acquire kind-like properties through repeated administrative use. They shape how populations are targeted, how data is gathered, and how political narratives are framed.

These examples suggest that regulatory kinds are characterized not by their semantic clarity, but by their institutional portability. A classification becomes regulatory when it is:

- **Recursively deployed** across multiple systems (health, law, policy)
- **Epistemically treated** as a kind — used to infer, predict, or justify action
- **Operationally portable** despite conceptual inconsistency

²⁸Angwin et al. (2016). “Machine Bias.” *ProPublica*.

²⁹Tremain, S. (2001). “On the Subject of Impairment.” *The Disability Studies Reader*.

- **Institutionally embedded** in forms, databases, protocols, or rules

This distinguishes regulatory kinds from other social kinds. Not all socially constructed categories are regulatory kinds. Some lack the cross-domain uptake. Others are too fluid to support institutional reasoning. The regulatory kind is a specific species within the broader genus of social kinds — marked by its function, not its content.

What unites these cases is not their metaphysical similarity, but their epistemic role: each serves as a surrogate for a more complex or inaccessible reality. Institutions act on these surrogates because they must. And over time, the simulation of kindhood becomes functionally indistinguishable from the real thing.

The regulatory kind framework thus invites a rethinking of classification itself. It shifts the focus from what a kind is to what a kind does — from ontology to epistemic labor. It reveals that in many domains, kindhood is not found but enacted. And it opens the door to analyzing how other categories — from “terror suspect” to “gender diverse” to “digital native” — may function as regulatory kinds in specific contexts.

5.4 Regulatory Kinds and the Ethics of Classification

To describe a classification as a regulatory kind is not to excuse its use. It is to analyze how it functions — and to ask what that function enables, obscures, or legitimates. Regulatory kinds stabilize action under uncertainty, but they also shape the epistemic and political conditions under which action becomes possible. They are not inert tools. They are structures of governance. And as such, they raise pressing normative questions.

First, there is the question of consent. Many regulatory kinds are assigned without deliberation, contestation, or appeal. An individual’s race, disability status, or risk tier may be recorded, inferred, or predicted without their knowledge — yet may shape the opportunities they receive, the scrutiny they attract, or the interventions they undergo. The simulation of kindhood can thus produce a form of epistemic capture: a person is acted upon as if they belong to a kind, even when that kind is conceptually unstable or empirically misleading.

Second, regulatory kinds can become instruments of inequality. By embedding social hierarchies into epistemic infrastructure, they may reproduce the very disparities they purport to manage. When race is treated as a medical risk factor, or poverty as a behavioral indicator, the classification itself may obscure structural causes, legitimize targeted surveillance, or rationalize uneven care. What appears as neutral stratification may be the institutionalization of bias.

Third, there is the problem of inertia. Once a regulatory kind is stabilized, it becomes difficult to revise or retire. Institutions built on recursive uptake resist change. This is

especially true when the classification supports compliance, funding, or automation. Even when critiques are known and acknowledged, the category remains because too much depends on its continued use. This stickiness presents a challenge for justice-oriented reform: the classifications most in need of scrutiny may be the least flexible in practice.

None of this entails a blanket rejection of regulatory kinds. Institutions must act, and some form of classification is unavoidable. The ethics of classification begins not with purity, but with reflexivity. What matters is recognizing the contingency of regulatory kinds — their constructedness, their effects, and their alternatives. This includes being attentive to who defines the kind, who is subjected to it, and who benefits from its use.

It also includes designing for corrigibility. Classifications should be contestable, revisable, and accountable. They should come with metadata: not just what the category is, but why it is used, how it was constructed, and what consequences attach to it. Regulatory kinds, if they are to serve ethically as well as functionally, must be made visible as instruments — not mistaken for facts.

Ultimately, the concept of regulatory kinds does not resolve the ethics of classification. It reframes it. It invites us to see categories not as mirrors of reality, but as architectures of inference and control. And it demands that we attend not only to what classifications are supposed to mean, but to what they do.

5.5 Simulated Kindhood, Real Consequences

This paper has introduced regulatory kinds as a metaphysical refinement for understanding how certain socially constructed classifications come to function epistemically across institutions. These are not merely social kinds in flux, nor looping kinds shaped by identification, nor ideologically saturated tools of critique. They are categories that behave like natural kinds within systems that need stable scaffolds for inference, coordination, and action.

The core insight is that kindhood can be simulated — not pretended or imposed, but enacted through recursive uptake. What distinguishes regulatory kinds is not their semantic clarity or causal unity, but their institutional portability. They achieve explanatory power not because they mirror nature, but because they are used as if they do. This is what gives them resilience, but also opacity.

The consequences are profound. Regulatory kinds shape the way we diagnose illness, allocate resources, assess risk, and measure equity. They determine who is legible within systems, and on what terms. And because they are often embedded in infrastructural routines — not in theory or policy alone — they tend to endure even after their conceptual basis is contested.

The framework developed here does not aim to replace existing models of classification, but to extend them. It brings into view a form of kindhood that is not discovered but

enacted, not found in nature but forged in bureaucratic dependence. It shows how institutions can stabilize uncertainty through epistemically thin, yet functionally indispensable, classifications.

This raises further questions. How might we detect regulatory kinds in practice — particularly when their reification is advanced? What normative criteria should govern their deployment or retirement? And what happens when regulatory kinds conflict, overlap, or fail? These questions extend beyond philosophy of science into the ethics of governance, the politics of data, and the sociology of institutions.

For now, the answer may be this: regulatory kinds help us understand why categories that seem unstable on paper remain indispensable in practice. Their kindhood is not an illusion. It is a consequence. And it is one we cannot afford to ignore.

Conclusion

The central argument of this paper has been that certain socially constructed classifications — despite lacking biological grounding or semantic coherence — come to function epistemically like natural kinds. They support inference, justify intervention, and structure institutional coordination. This is not because they reflect natural divisions in the world, but because they are treated as if they do. I have called such categories *regulatory kinds*, and described their epistemic function in terms of *simulated kindhood*.

To simulate kindhood is not merely to imitate its surface features. It is to occupy its institutional role. Regulatory kinds acquire their force through recursive uptake, infrastructural embedding, and functional necessity. They do not arise from discovery, but from dependence. And in doing so, they reshape the boundary between what is epistemically useful and what is ontologically justified.

Race, as the central case study, illustrates this dynamic vividly. It is unstable in meaning, heterogeneous in referents, and ethically fraught. Yet it endures in clinical guidelines, census schemas, forensic databases, and genomic studies. This endurance cannot be explained by its truth; it must be explained by its institutional function. Race is used to predict, allocate, and explain — not because it is a kind in the naturalistic sense, but because institutions have built themselves around the assumption that it behaves like one.

This insight complicates many familiar philosophical assumptions. It suggests that kindhood may be enacted rather than discovered, and that classification may be less a reflection of reality than a technique for governing under constraint. It also reveals a new source of epistemic inertia. When institutions simulate kindhood through classification, they generate outputs that appear to confirm the very distinctions they encode. The classification becomes its own evidence. In this way, regulatory kinds demonstrate a form of epistemic self-reference:

categories persist not by tracking stable referents, but by coordinating institutional action in ways that make them appear indispensable.

This presents both an analytic opportunity and an ethical challenge. Analytically, regulatory kinds offer a tool for understanding why some contested classifications endure across contexts, despite being the subject of sustained conceptual critique. They provide a vocabulary for describing classificatory resilience without reifying the categories themselves. Normatively, however, regulatory kinds raise urgent questions about institutional accountability, corrigibility, and the politics of epistemic infrastructure. If simulation can become reification — if institutional use can produce the appearance of truth — then scrutiny must shift from whether a classification is true to how it is used, by whom, and with what effects.

Ultimately, this paper advances a modest thesis with broad implications. It does not claim that all social kinds are regulatory kinds, or that regulatory kinds should be abandoned. It claims only that simulated kindhood is a central, under-theorized mechanism by which institutions manage complexity. To understand how classification governs, we must attend not just to what categories mean, but to how they function — and to what happens when epistemic scaffolds begin to shape the very world they were built to navigate.

References

- Angwin, J., Larson, J., Mattu, S., & Kirchner, L. (2016). “Machine Bias.” *ProPublica*. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Benjamin, R. (2019). *Race After Technology: Abolitionist Tools for the New Jim Code*. Polity Press.
- Boyd, R. (1991). “Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds.” *Philosophical Studies*, 61(1–2), 127–148.
- Bowker, G. C., & Star, S. L. (1999). *Sorting Things Out: Classification and Its Consequences*. MIT Press.
- Braun, L. (2014). *Breathing Race into the Machine*. University of Minnesota Press.
- Braun, L., Fausto-Sterling, A., Fullwiley, D., Hammonds, E., Nelson, A., Reverby, S., & Roberts, D. (2007). “Racial Categories in Medical Practice: How Useful Are They?” *PLOS Medicine*, 4(9), e271.
- Centers for Disease Control and Prevention (CDC). (2022). “Racial and Ethnic Disparities Continue in Pregnancy-Related Deaths.” <https://www.cdc.gov/mmwr/volumes/71/wr/mm7142a1.htm>
- DiMaggio, P. J., & Powell, W. W. (1983). “The Iron Cage Revisited.” *American Sociological Review*, 48(2), 147–160.
- Eberhardt, J. L. (2019). *Biased: Uncovering the Hidden Prejudice That Shapes What We See, Think, and Do*. Viking.
- Fujimura, J. H., & Rajagopalan, R. (2011). “Different Differences.” *Social Studies of Science*, 41(1), 5–30.
- Hacking, I. (1995). “The Looping Effects of Human Kinds.” In D. Sperber et al. (Eds.), *Causal Cognition* (pp. 351–394). Oxford University Press.
- Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford University Press.
- Hochschild, J., & Powell, B. (2008). “Racial Reorganization and the United States Census 1850–1930.” *Studies in American Political Development*, 22(1), 59–96.

- Kahn, J. (2013). *Race in a Bottle*. Columbia University Press.
- Kuzawa, C. W., & Sweet, E. (2009). “Epigenetics and the Embodiment of Race.” *American Journal of Human Biology*, 21(1), 2–15.
- Morning, A. (2011). *The Nature of Race*. University of California Press.
- Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). “Dissecting Racial Bias in an Algorithm.” *Science*, 366(6464), 447–453.
- Office of Management and Budget (OMB). (1997). “Revisions to the Standards for the Classification of Federal Data on Race and Ethnicity.” *Federal Register*, 62(210), 58781–58790.
- Roberts, D. (2011). *Fatal Invention*. The New Press.
- Sauer, N. J. (1992). “Forensic Anthropology and the Concept of Race.” *American Anthropologist*, 94(1), 107–115.
- Spencer, Q. (2014). “A Radical Solution to the Race Problem.” *Philosophy of Science*, 81(5), 1025–1038.
- Spencer, Q. (2019). “What Is Race?” In S. Haslanger & Á. Ásta (Eds.), *The Routledge Handbook of Social Ontology*. Rowman & Littlefield.
- Tremain, S. (2001). “On the Subject of Impairment.” In G. L. Albrecht et al. (Eds.), *The Disability Studies Reader* (pp. 282–291). Routledge.