


COGNITIO

Revista de Filosofia
Centro de Estudos de Pragmatismo

São Paulo, v. 26, n. 1, p. 1-15, jan.-dez. 2025
e-ISSN: 2316-5278

 <https://doi.org/10.23925/2316-5278.2025v26i1:e70312>

From indeterminacy to normativity: Quine, meaning, and teleosemantics

Da indeterminação à normatividade: Quine, significado e teleossemântica

Lucas Ribeiro Vollet*
luvollet@gmail.com

Abstract: The problem of meaning indeterminacy has long challenged philosophers, particularly at the intersection of language, cognition, and epistemology. While teleosemantics offers a promising naturalistic framework by grounding meaning in the evolutionary functions of biological systems, it faces a persistent challenge: how to account for the normative dimensions of meaning — those that make interpretations *justifiable* rather than merely *adaptive*. This paper addresses that challenge by integrating insights from Quine’s critique of analyticity and his holistic epistemology with teleosemantic theories developed by Millikan and Papineau. We argue that meaning stabilizes not through formal structures alone, nor through mere causal success, but via historically embedded epistemic practices that reflect a capacity for reflexive self-correction. This reflexivity, we suggest, enables a form of conceptual self-consciousness grounded in our biological background rather than in metaphysical assumptions — thus securing a continuity between empirical inquiry and normative justification. By framing normativity as an emergent property of coordinated, self-aware epistemic communities, the paper challenges reductionist accounts of meaning and advocates for a refined naturalism — one that preserves the critical tension between descriptive success and the evolving standards of rational endorsement.

Keywords: Meaning Indeterminacy, Normativity, Quine, Teleosemantics

Recebido em: 11/02/2025.

Aprovado em: 25/06/2025.

Publicado em: 03/12/2025.

Resumo: *O problema da indeterminação do significado há muito desafia os filósofos, particularmente na interseção entre linguagem, cognição e epistemologia. Embora a teleossemântica ofereça um quadro naturalista promissor ao fundamentar o significado nas funções evolutivas dos sistemas biológicos, ela enfrenta um desafio persistente: como explicar as dimensões normativas do significado — aquelas que tornam as interpretações justificáveis e não apenas adaptativas? Este artigo aborda esse desafio integrando a crítica da analiticidade a epistemologia holista quiniiana com as teorias teleossemânticas desenvolvidas por Millikan e Papineau. Argumentaremos que o significado não se estabiliza apenas por meio de estruturas formais, tampouco pelo mero êxito causal, mas sim através de procedimentos epistêmicos historicamente enraizados que refletem uma capacidade reflexiva de autocorreção. Essa reflexividade, sugerimos, possibilita uma forma de autoconsciência conceitual ancorada em nosso arcabouço biológico, e não em pressupostos metafísicos — assegurando, assim, uma continuidade entre a experiência e a justificação normativa. Ao enquadrar a normatividade como uma propriedade emergente de comunidades epistêmicas coordenadas e autoconscientes, o artigo desafia conceituações reducionistas do significado e defende um naturalismo refinado — capaz de preservar a tensão crítica entre o êxito descritivo e os padrões evolutivos de endosso racional.*

Palavras-chave: Indeterminação do significado. Normatividade. Quine. Teleossemântica.



Artigo está licenciado sob forma de uma licença
Creative Commons Atribuição 4.0 Internacional.

1 Introduction

This article explores the transition from Quine’s critique of intensional objects—such as descriptions, possible worlds, and counterfactuals—

* Universidade Federal de Santa Catarina.

to contemporary attempts to resolve the problem of meaning indeterminacy through naturalistic approaches. Quine's skepticism toward intensionality catalyzed a broader philosophical shift, one that favored empirical and functionalist accounts of meaning over abstract semantic structures. Yet while Quine effectively undermined the notion of fixed meanings and highlighted the underdetermination of translation, he left open a fundamental question: how can content ever be sufficiently specified to support communication, scientific theorizing, or conceptual progress?

If we take the naturalist framework seriously, indeterminacy should not be seen as a flaw to be eliminated, but rather as a generative feature—an open space in which shared assumptions consolidate through practical engagement and interpretative negotiation. This view is compatible with the core of Quine's semantic skepticism: the flexible boundaries of core meaning, or analyticity (Quine, 1953), and the fundamentally evolving, socially embedded nature of meaning formation (Quine, 1960). Rather than undermining objectivity, such indeterminacy highlights the contingent yet corrigible scaffolding upon which linguistic practices and conceptual schemes are built.

However, this interpretative openness has its limits. Even the most sophisticated versions of teleosemantics, while promising in their grounding of meaning in the evolutionary functions of biological systems, often stop short of explaining how practical success translates into normative authority. What remains unaddressed is the synthesis of successful behavior into epistemically justified, normatively binding representations.

Our aim, then, is to interrogate whether naturalistic models—particularly teleosemantic ones—can bridge this gap. Can the notion of *success*—whether defined through reward-punishment mechanisms, the stabilization of reliable associations, or inductive generalizations—truly ground the normative force of meaning? Or does it risk reducing normativity to a mere biological contingency or arbitrary social construct?

These questions matter because they target the foundational tension in theories of meaning: the divide between behavioral regularity and epistemic justification. If meaning were solely a function of communicative efficacy, then any stable pattern could, in principle, qualify as meaningful. But this overlooks the deeper normative concern—*what makes an interpretation not only successful but right?*

Teleosemantics provides a powerful answer to Quine's challenge by locating semantic content in the causal-historical functions of representational systems. But to move beyond Quine's underdetermined isolation of reference, it must also supply a framework for typifying representational mechanisms—one that explains, for example, why the representation of a frog differs from that of a human, or how scientific representations achieve a level of fine-grained nuance and internal coherence that superstition lacks. Without such differentiation, naturalistic accounts risk flattening the semantic landscape, ignoring the graded, structured, and epistemically accountable dimensions of content.

This article contends that the true strength of teleosemantics lies not in bypassing normativity, but in showing how it *emerges*—reflexively and historically—from within biologically grounded representational systems. Instead of relying on formal rules or metaphysical guarantees, we argue that normativity arises through belief revision and epistemic refinement over time. In this respect, Quine's continuity thesis—when extended and critically elaborated—offers the beginnings of an answer: a path from descriptive success to normative justification, grounded not in metaphysics, but in the evolving interplay of naturalistic inquiry and conceptual self-correction.

Moreover, the article suggests that normativity—often dismissed as a metaphysical leftover or reduced to social convention—may be better understood as an emergent feature of reflexive, historically embedded inquiry. This perspective engages an ongoing challenge within cognitive science: how to incorporate normative reasoning into models grounded in embodiment and functional success. By emphasizing conceptual self-correction rather than behavioral convergence alone, we move toward a framework that not only accounts for how systems coordinate and adapt but also explains how they

distinguish between interpretations that are equally effective in practice yet carry different epistemic commitments or justificatory weight.

2 Quine's indeterminacy: from sensory input to semantic fragility

Quine's thesis on the indeterminacy of translation, mainly articulated in *Word and Object* (1960/2013), remains a central challenge in the philosophy of language. According to Quine: "manuals for translating one language into another can be set up in divergent ways, all compatible with the totality of speech dispositions, yet incompatible with one another" (Quine, 2013, p. 24). In short, there is no fact of the matter that can determine which among multiple translation manuals is the "correct" one; empirical data underdetermines semantic assignments. This problem is particularly acute when meanings are assumed to be recoverable from behavior or sensory stimuli alone.

In *Word and Object* (1960), Quine introduced the notion of stimulus meaning as a way to naturalize semantics, grounding linguistic behavior in observable dispositions. Yet he was clear that this notion was never meant to function as a traditional semantic category. In *The Roots of Reference* (1974), he reframes stimulus meaning as an epistemological tool, useful for tracing how language acquisition occurs via conditioning and social reinforcement rather than introspection or reference to inner ideas. Observation sentences are characterized not by their introspective basis but by the possibility of intersubjective agreement: "Observation sentences are the gateway to language" (Quine, 1974, p. 41). Rather than positing mental entities, Quine insists that our grasp of meaning begins in shared contexts of perceptual similarity and assent. Though often misunderstood, Quine did not renounce the idea of stimulus meaning but refined its scope, insisting that it remain anchored in the subject's sensory triggers rather than in shared distal objects—a stance reaffirmed in his late essays such as *Progress on Two Fronts* (1996) and *Pre-Established Harmony* (published posthumously in Janssen-Lauret & Kemp, eds., *Quine and His Place in History*, 2016).

Holism, not *stimuli* alone, became the cornerstone of Quine's mature epistemology. As he argued in *Two Dogmas of Empiricism* (1953), statements face the tribunal of experience not in isolation but as part of a network of beliefs. This network, constantly revised and restructured, underlies our conceptual practices, including science. However, Quine kept his semantic skepticism. He warns against overestimating how well intrasubjective stimulus synonymy holds up under the pressure of collateral information. As he writes in *Word and Object*:

Our success with 'bachelor' and 'unmarried man' has been sufficient, despite the impasse at second intention, to tempt us to overestimate how well intrasubjective stimulus synonymy withstands collateral information. By way of corrective, consider the Himalayan explorer who has learned to apply 'Everest' to a distant mountain seen from Tibet and 'Gaurisanker' to one seen from Nepal. As occasion sentences these words have mutually exclusive stimulus meanings for him until his explorations reveal, to the surprise of all concerned, that the peaks are identical. His discovery is painfully empirical, not lexicographic; nevertheless, the stimulus meanings of 'Everest' and 'Gaurisanker' coincide for him thenceforward. (Quine, 2013, p. 44).

This means that the apparent agreement in our dispositions to assent and dissent cannot serve as a justificatory bridge to a richer dimension of meaning—one populated by possible worlds, private imagery, or other hypothetical cognitive preconditions. Because Quine never relinquished his skepticism, the absence of a richer semantic dimension leaves unresolved a host of problems concerning the stability of the very practices that rely on meaning as a tool—practices that are not merely idle exchanges, but often directly tied to our survival. He left open the question of how stable meanings emerge and persist despite the indeterminacy of translation. Teleosemantics enters this landscape not in opposition to Quine but as

an extension of his naturalistic thesis. Where Quine naturalized epistemology, teleosemantics attempts to naturalize semantics—explaining the stabilization of meaning through the evolutionary functions of representational systems. By tracing how adaptive pressures select for reliable patterns of environmental representation, teleosemantics reframes the question of reference: not “which translation is correct?” but “which representational function has historically contributed to an organism’s success?”

Yet even this move faces a challenge. If semantic content is grounded in evolutionary success, what distinguishes cognitively meaningful distinctions—such as scientific taxonomies—from arbitrary or superstitious ones? Bridging this gap requires more than adaptive history; it requires a normative framework through which representational commitments are justified, refined, and critically evaluated. This is the space where our article intervenes: arguing that teleosemantics, to succeed, must incorporate a layer of reflexive self-awareness capable of transforming practical success into epistemic legitimacy.

3 From Semantic arbitrariness to functional normativity: teleosemantics and the evolutionary shaping of meaning

Classes are cheap. Categories or classifications are readily available and trivially generated—one can, as Nelson Goodman famously demonstrated, carve nature at its joints or at its scars, depending on one’s choice of predicates (Goodman, 1978). The syntactic machinery of language permits a virtually infinite recombination of features and criteria, meaning that from a formal standpoint, any object can belong to as many classes as there are conceivable predicates.

But how does one cross the bridge from semantic arbitrariness to genuine normativity? In other words, how can we distinguish between classifications that simply *are* and those that *ought* to constrain cognition and behavior? And crucially, how can this be done without resorting to metaphysical or nomological residues—the kind of unverifiable, quasi-mystical essences that J.J.C. Smart famously warned against in his 1959 critique of non-naturalistic mind-talk (*Sensations and Brain Processes*)? This concern resurfaces in the eliminativist arguments advanced by Paul Churchland (1981), who likened our folk-psychological concepts to obsolete scientific theories—useful fictions lacking explanatory power—and by Patricia Churchland (1986), who maintained that any meaningful discussion of normativity in cognition must be firmly rooted in neurobiology. More recently, Ruth Millikan’s teleosemantic framework (1984, 1993) offers a naturalized account of meaning and content without invoking metaphysical fixities—showing how biological functions and evolutionary history can ground normative distinctions without slipping into essentialism.

Teleosemantics offers a compelling and naturalistic response to this dilemma. It ties semantic correctness not to abstract properties or transcendent truths, but to historical biological functions. On this view, the “right” class is the one selected for in the evolutionary history of a system—because it promoted survival, reproductive success, or reliable environmental responsiveness. Conversely, the “wrong” class is one that leads to systematically maladaptive outcomes.

This approach reframes normativity not as a mysterious, mind-transcendent force, but as a biologically anchored consequence of functional success. Not all classes are equal: some support survival, others lead to error. In this way, teleosemantics crosses the bridge from arbitrariness to normativity without invoking ghostly metaphysical predicates—only historical, causal, and adaptive patterns.

4 Fixing content without metaphysics: teleosemantics beyond Quinean indeterminacy

Teleosemantics posits that the content of a representation is determined by the evolutionary function it serves for an organism, thereby linking meaning to naturalistic, causal processes. On this view, a biological

system's survival pressures play a decisive role in shaping what its mental representations are "about." For instance, a frog's ability to catch flies relies on its perceptual system recognizing certain moving stimuli—such as small, dark, rapidly moving dots—as flies. The content of this mental representation is said to be shaped by the frog's adaptive function, as its survival depends on accurately identifying prey. The perceptual system, through evolutionary pressures, becomes specialized to detect specific features of its environment, and it is this evolutionary history that fixes the content of its representations.

The teleosemantic thesis can be seen as a natural development from Quine's confrontation with meaning indeterminacy—especially in the domains of translation and modal discourse, where competing manuals of interpretation resist convergence. In Quinean terms, this is a consequence of the underdetermination of theory by evidence: empirical data underdetermine which conceptual scheme is "correct," because multiple incompatible schemes can accommodate the same sensory input. But while Quine exposes this ambiguity, he doesn't resolve it—he suspends the notion of fixed meaning in favor of a pragmatic web of beliefs. Teleosemantics, by contrast, reintroduces a kind of constraint: not through logical necessity, but through evolved functionality.

What Quine describes as a "manual of translation" finds a natural analogue in teleosemantics as a set of evolved, causally grounded correlations between signs and environmental features. When these correlations are recurrent and functionally reliable, they approximate what Quine ultimately withheld: a principled, non-arbitrary mapping between language and world. Teleosemantics thus extends—and radicalizes—Quine's insight. Where Quine maintains that meaning emerges from a sentence's place within an evolving web of belief, teleosemantics claims that representations earn their content through a history of adaptive success.

In this framework, what matters is not mere extensional agreement between interpretations, but whether they play comparable functional roles within systems of inference, learning, and adaptation. Interpretations that coincide in reference may still diverge radically in how they guide action, shape expectations, or integrate into broader networks of belief. The key distinction lies in historical performance: an interpretation earns its place not through conceptual fiat, but through its sustained success in navigating the world. What teleosemantics emphasizes—building on Quine's holism—is that representational content is forged through patterns of correction, responsiveness, and resilience over time. Meaning, in this sense, is not fixed by definition but stabilized through repeated interaction, where selective pressures reward those mappings that prove most conducive to survival.

The teleosemantic thesis presents itself, in this context, as a rather elegant strategy. By grounding meaning in the natural history of organisms—specifically in their evolutionary functions and success conditions—teleosemantics introduces an external criterion for fixing content that does not rely on unverifiable metaphysical assumptions. It thus offers a way to preserve the empirical tractability of semantic claims, aligning with the efforts to tie meaning to systematic, observable, and intersubjectively controllable procedures.

In other words, teleosemantics seeks to explain the shift from mere content computation and extensional enumeration to the strong and normative identification of 'content.' The normativity that emerges is not an arbitrary constraint, but a structural necessity shaped by the organism's threshold for sustainable error—its capacity to tolerate deviations before those deviations become maladaptive for survival. In this sense, the mystery surrounding normativity diminishes when it is understood not as an abstract imposition but as the outcome of specific evolutionary pressures that delineate the permissible range of interpretative flexibility and classificatory efficiency. For instance, the human ability to recognize faces is not a random or arbitrary feature; it has evolved as a critical adaptation, where deviations in recognizing faces (e.g., misidentifying people) could lead to social and survival disadvantages.

In *Language, Thought, and Other Biological Categories* (1987), Millikan emphasizes that Quine's problem of indeterminacy applies to systems of arbitrary signs, highlighting that the emergence of genuine intentionality and representational content is dependent on conditions of increasing complexity.

However, these conditions can only take on a seemingly simple, common, and universal form if they are understood from the perspective of prolonged exposure to evolutionary time. In other words, what appears to be a straightforward mapping of signs to meanings is, in reality, the outcome of a historical process in which cognitive structures are gradually shaped by adaptive pressures.

5 Swampman and the problem of empty content: testing the boundaries of teleosemantics

Section 2 showed how teleosemantics naturalizes meaning by grounding it in evolutionary function, thereby addressing the indeterminacy problem without resorting to metaphysical entities or analytic dogmas. Yet this account raises a deeper issue: can success in evolutionary terms alone fully explain why a particular representation *ought* to count as meaningful? Is the selective history of a system sufficient to generate normativity, or does this merely displace the problem onto biology? In other words, does evolutionary success confer meaning, or does it only confer reliability? And further: is reliability—being a system whose actions are consistently predictable across contexts—sufficient to endow those actions with genuine content? Or does content require more than behavioral regularity—something like a history of interpretative embedding or epistemic accountability?

These questions become especially pressing when we consider borderline cases—scenarios that meet physical or behavioral criteria, yet intuitively lack intentional content. Donald Davidson’s Swampman thought experiment dramatizes this tension: a being that behaves identically to a person, but without any causal history, seems to lack representational content altogether. Davidson (1987) asks us to imagine that, after being struck by lightning, he is instantly vaporized, and in the same moment, by sheer chance, a molecule-for-molecule duplicate of him is spontaneously formed in a nearby swamp. Though this *Swampman* would behave indistinguishably from Davidson—writing the same papers, recognizing the same friends, and appearing to have identical thoughts—Davidson insists that it could not possess genuine intentional states. Lacking a causal history of interaction with the world, the *Swampman* could not have acquired the referential connections that underpin meaningful thought and language.

Millikan’s teleosemantic account provides a naturalistic grounding for this intuition, explaining why content must be tied to historical selection rather than mere physical structure. For her, intentionality is not just a matter of having the right physical arrangement but of having a cognitive system shaped by evolutionary pressures. She emphasizes that the *Swampman*, despite its outward indistinguishability from Davidson, would lack intentional states because it has not undergone the selection process that defines proper functions: “only in virtue of one’s evolutionary history do one’s mental intentional states have proper functions” (Millikan, 1987, p. 93).

The key point here is that, despite its ability to behave as if it understands and thinks, the *Swampman*’s *propositions* would be akin to a randomly generated string of letters that coincidentally forms an intelligible sentence—it lacks any grounding in the processes that normally give such a sentence meaning. Millikan presses this issue further: “I see no way of unpacking the should in this sort of context by reference to present structure or disposition. How could any ‘should’s’ or ‘supposed to’s’ be applied to inner arrangements of our newly arrived, randomly created double?” (1987, p. 94). As Millikan provocatively concludes, Quine’s indeterminacy thesis would apply to this being’s utterances “with a vengeance never envisioned by Quine” (1987, p. 93), highlighting the radical extent to which intentional states are historically constituted rather than merely extensionally instantiated.

Millikan’s broader project extends beyond this thought experiment, aiming to naturalize intentionality through what she calls *proper functions*—biological purposes shaped by natural selection. Unlike traditional functionalist theories, which focus on structural role alone, her account insists that representational content depends on the evolutionary pressures that have historically reinforced certain

cognitive and perceptual mappings over others. For Millikan, meaning is not just a product of causal interactions but of a history of selection, where successful representation-tracking is tied to survival and reproduction. This view has profound implications, as it challenges any attempt to detach semantics from biology, grounding meaning firmly in the adaptive landscape that has shaped our cognitive capacities.

Thus, meaning is not a mere abstraction nor a product of arbitrary convention, but a reflection of how cognitive systems, linguistic practices, and social structures co-evolve. The challenge, then, is recognizing that even our most “objective” classifications are embedded within this evolutionary and epistemic scaffolding—one that not only enables meaning but also constrains and stabilizes it within historically contingent yet normatively binding frameworks.

David Papineau (2022) defends the traditional aetiological (historical) teleosemantic approach, arguing that selectional histories are crucial to understanding representation. Instead of relying on intrinsic properties like molecular structures to explain representational content, Papineau asserts that it is the relational historical properties that determine what counts as representation.

Papineau draws an analogy between representational systems and scientific classifications such as water (H₂O). Just as the real essence of water lies in its molecular composition, the real essence of representation, according to Papineau, lies in its selectional history. He argues that representational systems constitute a natural kind, unified by their evolutionary origins and shaped by the selective pressures they have faced across generations. This shared selective provenance, he claims, explains why different species develop similar representational structures despite distinct evolutionary paths. Papineau views this selective history as essential to understanding representation, countering thought experiments like the Swampman scenario, which posit that representation could arise without any historical basis. He argues that representational systems are a natural kind, unified by the evolutionary origins and selective pressures they faced:

As with other functional kinds, we can thus view these common selective pressures as the super-explanatory property of the kind. Their common selective provenance is the property that explains why all instances of representation share a range of other properties. Given this, it is natural to view this selective provenance as the providing the “essence” of representational systems. (Papineau, 2022, p. 14).

Papineau thinks that the essence of representational systems is their shared selective provenance, which explains why they display similar structures across species. For him, this historical dimension is critical in rejecting the idea that mental states can be reduced to intrinsic, present-day properties alone. He argues that representational content is fundamentally tied to the functions for which cognitive states were selected. As he puts it, “Belief is not an implicit description of whatever brain state satisfies some everyday descriptive role, with the implication that some non-selected brain state might satisfy that role in a possible swampman scenario” (PAPINEAU, 2022, p. 17).

What these accounts (Millikan and Papineau) make clear is that behavioral reliability, no matter how consistent, is not sufficient to establish representational content. Without a history of selection or functional differentiation, such regularity risks being mistaken for meaning where there is none. Evolutionary success explains why certain mappings persist, but it does not fully explain why they *ought* to be treated as meaningful in epistemic or normative terms. In this light, the Swampman thought experiment is not merely a counterexample—it reveals a deeper tension: even when an agent’s actions are predictable, the absence of an anchoring historical function deprives those actions of intentional significance.

What remains open for the next chapter, then, is whether teleosemantics can account not only for the stability of content but for its legitimacy. To address this, we must consider whether biological function alone is sufficient to ground normativity, or whether a second layer—epistemic, reflective, and socially negotiated—is required to transform functional success into genuine meaning.

6 From pragmatic success to epistemic normativity: bridging teleosemantics and Quinean naturalism

Papineau critiques non-aetiological theories of teleosemantics, which claim that representational content can be understood without reference to historical selection processes. He argues that these theories fail to account for the natural kind essence of representation, as they overlook the importance of selectional history in grounding representational norms. This view aligns with Quine's recognition that knowledge is inherently revisable and that the boundary between analytic and synthetic statements is not fixed but negotiated within the evolving fabric of our theoretical commitments. For both Quine and Papineau, meaning is not grounded in immutable distinctions but in the dynamic interplay between belief systems and empirical constraints. The analytic/synthetic divide, under this perspective, is not a logical given but a historically contingent artifact of how particular frameworks manage to organize experience more coherently than their rivals. When a specific interpretation of meaning becomes dominant—when it takes on the appearance of necessity—it does so not through metaphysical privilege, but because it has proven more resilient, more predictive, and less prone to failure than alternative schemes. Teleosemantics extends this insight by situating that process within the pressures of evolutionary history: it is the selective filtering of representational strategies that explains why certain meanings stabilize and gain normative force. In this sense, meaning does not emerge from abstract stipulation but from the success of historically grounded systems in navigating a world that pushes back.

Indeed, if we take Quine's revisability thesis seriously, we must accept that even our most secure categories of thought — whether analytic or synthetic — are provisional. They are outcomes of evolutionary histories, social negotiations, and practical constraints, not timeless truths. This allows us to make sense of why representational frameworks between species might occasionally overlap, without positing a shared conceptual schema. It provides a compelling rationale for understanding why we converge to a certain extent with some species and not with others. If something explains why dogs, for instance, do not seem to categorize their interactions with humans in terms of political or institutional roles—whereas humans routinely interpret relationships through categories like kinship, authority, or labor—it cannot be solely a matter of neural architecture or immediate cognitive capacity. Rather, it points to a deeper historical divergence: the history of each species' interaction with its environment, its selective pressures, and its ecological niche. These distinct histories give rise to distinct representational frameworks. To be sure, overlaps do occur—dogs can be trained to recognize commands or respond to boundaries—but if such convergence happens, it too demands explanation. And if not purely by present function, then perhaps by traces of a shared evolutionary history that afford moments of partial alignment.

After all, a dog can be trained to recognize a boundary, much like a human recognizes a national border. But does this behavioral convergence indicate a shared conceptual framework, or merely a pragmatic alignment born of functional conditioning? The dog may respond appropriately to cues in its environment, but it does not operate within the same inferential network that allows a human to distinguish between political sovereignty and territorial demarcation. What we observe, then, is not necessarily shared meaning, but a *margin of tolerance*—a limited range within which different cognitive systems can produce similar responses because they face similar practical demands. In such cases, coordination does not imply that the systems interpret the world in the same way; it only shows that their behaviors align under specific conditions. Meaning, on this view, is not something inherent in a representation itself, but rather the result of an evolutionary convergence—where different organisms, shaped by similar environmental pressures, develop functionally compatible responses. The most basic form of convergence between living beings is simply the fact that they are alive and must survive in roughly similar worlds. This shared condition is enough to ensure that one system's successful behavior often correlates with another's—though never perfectly, since survival permits a certain degree of

variation. Thus, the margin of error is bounded, not because the systems share conceptual content, but because both are subject to the same biological constraints.

This notion becomes especially relevant when extended to the Swampman: if he interacts with us as fluently as our dogs do—responding to instructions, navigating spatial relations, producing meaningful utterances—this fluency may conceal a deeper absence. When the Swampman says “Brazil is a sovereign country,” he may not *mean* what we mean, in the rich, historically embedded sense of that claim. Yet when he says “this is to the right of that,” he converges with us in a basic functional task. But this convergence, like the dog’s training, may not suffice to establish true conceptual sharing—it may reflect only a structurally consistent behavior within a minimal margin of compatibility.

This directly engages with Quine’s problem of indeterminacy by reframing the *margin of tolerance* as a bounded form of interpretative flexibility—a degree of indeterminacy that, while real, does not result in communicative breakdown or unbridgeable divergence. In other words, systems can differ in how they represent the world, yet still coordinate effectively within a shared zone of functional overlap. We may disagree, for example, on the precise conceptual boundaries of a term like “justice” or “nationhood,” but only up to a certain point; beyond that point, mutual understanding begins to falter. This bounded indeterminacy is not a flaw but a feature—it reflects the minimal convergence required for successful interaction in a world where meaning is shaped by use, adaptation, and correction rather than by fixed conceptual essences.

Rather than viewing meaning as entirely underdetermined, this perspective highlights how the content of representations is not fixed solely by their present structure but is historically shaped by the selective pressures that have guided their development. In this way, indeterminacy is not a radical openness but a constrained flexibility, where interpretative stability emerges not from rigid definitions but from the evolutionary and functional demands that have sculpted cognitive architectures over time.

This means that what determines these representational structures is not a grasp of logical necessity, inductive reasoning or an innate modal grasp, but the practical need to ensure reliable coordination between cognitive systems and their environments. Papineau writes, “Without historical selection, representational states would lack the relational properties that give them their explanatory power” (Papineau, 2022, p. 514). Without a historical basis, the content of a representation would remain ambiguous, as alternative interpretations could always be posited without any grounding in the organism’s evolutionary past. To sum up, we saw Papineau’s historical teleosemantic approach demonstrates that the causal mechanisms underlying representation are not reducible to mechanical determinism. This thesis aligns directly with key consequences of Quine’s philosophy: the rejection of any fixed analytic-synthetic distinction and the abandonment of modal notions as foundational for understanding empirical knowledge. For both Papineau and Quine, normativity arises not from inherent structures or immutable truths, but from practical interactions shaped by historical contingencies. The stability of meaning, rather than being a given, is achieved through the convergence of selective pressures that guide organisms toward predictable, functional interactions with their environments. Normativity, then, is not *a priori* but emerges from the historical success of certain representational practices over others. The only relevant trace of normativity for an empirical theory, therefore, lies beyond any static framework of analyticity or inherent modal insight. Instead, it is found in the gradual historical stabilization of representational systems through patterns of practical success—patterns that organisms are pressured to adapt to to survive and thrive. This historical stabilization is what ensures that organisms achieve stable interactions with the world. By grounding content in selective history, teleosemantics provides a way to move beyond the limitations of extensional sameness and toward a robust account of why certain interpretations hold normative force. However, the epistemic level of this normative selection is not considered. Right and wrong are not determined by any epistemic selection criterion but rather as a collateral outcome of the organism’s survival.

7 The normativity gap: from evolutionary success to epistemic justification

As we've now seen, teleosemantics provides a powerful framework for explaining how representational content can emerge from the functional success of cognitive systems. By grounding meaning in evolutionary history and adaptive reliability, it offers a naturalistic account of how semantic stability arises over time. This explains how organisms—including humans—come to coordinate representations and achieve communicative success. But a deeper challenge remains: how do we move from such descriptive facts about how representations function to the normative structures that determine what they *ought* to mean?

In other words, it is one thing to say that a population has historically converged on certain associations—say, that mermaids are imagined as pale-skinned—but quite another to explain why this association should carry any normative force, especially when the concept itself is fictional and fluid. Merely noting that such associations have stabilized tells us little about their legitimacy. Do they reflect arbitrary cultural habits? Do they conceal ideological operations? Or is there some rational structure that justifies their persistence? These are not peripheral concerns—they strike at the heart of what meaning is *for* us as reflective beings.

Teleosemantics answers part of this challenge. It shows how meaning can stabilize across a population through shared selective pressures and patterns of success in environmental interaction. This is sufficient to explain why we *do* coordinate meaning, why misunderstandings are bounded, and why certain expressions endure. But it is not enough to explain why one interpretation is *right* while another is *misguided*, or why we ought to revise certain frameworks in light of new insight. That is, biological stabilization may explain *how* communication works, but not *why* it should work this way—and not otherwise. It offers an account of convergence, but not of justification.

The issue of meaning stabilization extends beyond the mere effectiveness of human communication. While success in interpretation matters, it does not, by itself, address deeper concerns about the nature of that success—whether it is genuinely compelling, sustainable, and anchored in both a stable sociological framework and a biological rationale, rather than being shaped by ideological expediency. Is it grounded in a rationality that asserts itself normatively, or is it merely the result of an improvised, contingent adjustment? These questions, often dismissed as both radical and clichéd within philosophy, expose a real problem: there are inherent limits to how we establish criteria for successfully determining meaning. Eventually, we reach a threshold where the very notion of success becomes unstable, inviting precisely the kind of philosophical scrutiny that transforms Quine's semantic skepticism from an abstract puzzle into a source of genuine and unavoidable unease.

8 Naturalism with a reflexive turn: continuity and normativity without metaphysics

The preceding chapter revealed a structural tension at the heart of naturalistic accounts of meaning. While teleosemantics succeeds in explaining how content stabilizes through evolutionary success, it leaves unanswered the deeper question of why certain meanings should be preserved, revised, or replaced. In other words, the model explains how meaning is sustained but not why it ought to matter. The case of the Swampman illustrates this limitation vividly: though behaviorally indistinguishable from a fully developed person and capable of *successful* coordination, the Swampman lacks the historical and functional grounding necessary for its representations to carry normative weight. Its apparent *success* in mirroring meaningful behavior conceals a fundamental absence—namely, the lack of an epistemic lineage through which its commitments could be justified, revised, or held accountable. This highlights the insufficiency of mere functional convergence and underscores the need for a layer of reflexive evaluation through which success acquires normative significance.

But this raises a crucial question: if evolutionary success alone fixes content, does it truly place us in a better position than the Swampman? Is our meaning-making genuinely normative, or merely more stable? We argue that evolutionary success, while necessary, is not sufficient. What makes our representational practices more than sophisticated mimicry is not just their survival value, but their participation in a historically embedded, socially mediated, and critically self-correcting epistemic process. Only through this reflexive dimension can meaning acquire justificatory force—becoming not just a product of what worked, but a commitment to what can be evaluated, challenged, and improved.

In other words, teleosemantics offers the tools to track representational success, but lacks the internal architecture for evaluating which success-tracking mechanisms are epistemically justified—why some patterns deserve to be treated as knowledge, while others remain mere coordination artifacts. If we stop at mere functional convergence, we risk reducing semantic normativity to nothing more than a pattern of coordination with no intrinsic justificatory force. The challenge, then, is to remain committed to a naturalist framework without collapsing into either metaphysical speculation or empirical descriptivism.

This article argues that a more refined form of naturalism—one deeply informed by Quine’s continuity thesis, *the idea that epistemology is a chapter of psychology, and that philosophy is continuous with science* (Quine, 1969; 1975)—offers a compelling alternative to both foundationalist metaphysics and epistemic skepticism. Rather than seeking to reintroduce immutable first principles, this approach grounds normativity in a historically situated, epistemically self-aware process of conceptual evolution. On this view, science and philosophy are not sharply divided enterprises but part of a continuous fabric of inquiry—one in which necessary truths emerge not from *a priori* fiat, but as revisable outcomes of a dynamic, self-correcting scientific consciousness shaped by philosophical reflection.

While inspired by Quine’s rejection of any sharp boundary between empirical and conceptual domains, the proposal deliberately ventures into terrains Quine himself chose not to explore, particularly in granting a more constructive role to normativity and conceptual refinement within the scientific enterprise. In doing so, it extends Quinean naturalism beyond its original contours, supporting a model in which what counts as a necessary truth is not anchored in fixed analytic categories, but emerges through historically situated practices of rational inquiry and conceptual refinement.

It rests on five key ideas:

1. Continuity Without Reductionism.

Quine’s continuity thesis dissolves the sharp boundary between epistemology and empirical science. However, this continuity need not entail reductionism. The epistemic practices that emerge within science—practices of justification, revision, and conceptual evaluation—can retain their normative force without invoking transcendental foundations. Meaning is shaped by our interaction with the world, but also by the structures of inquiry we develop to make sense of that interaction.

2. Epistemic Reflexivity as a Natural Product

A truly robust naturalism must account for the fact that science does not just produce results—it produces standards for evaluating those results. This reflexive dimension is not an external philosophical add-on, but a constitutive feature of scientific practice. As natural science evolves, it generates its own criteria of success, internal standards of coherence, and forms of methodological self-correction. These are the seeds of normativity, growing from within the scientific process itself¹

1 On this point, see also adjacent accounts of conceptual development and justification in Kuhn’s *The Structure of Scientific Revolutions* (1962), Brandom’s *Making It Explicit* (1994), and Habermas’s theory of communicative rationality (*The Theory of Communicative Action*, 1981). While distinct in their premises, these frameworks similarly foreground the historical and socially embedded dimensions of normative emergence.

3. Justification Through Conceptual Co-evolution

Justification of meaning-claims does not lie in correspondence with metaphysical entities, but in their role within an evolving network of beliefs. Concepts are not fixed by biology alone; they are refined through iterative cycles of hypothesis, feedback, and reinterpretation. This model echoes Quine's holism: beliefs gain epistemic weight through their capacity to cohere with and support the wider system. Normativity thus emerges not from an absolute ground, but from the dynamics of conceptual interdependence and long-term refinement within a shared epistemic community.

4. Resistance to Dogmatism and Ideological Fixity

Without this reflexive layer, stabilization risks becoming ideological naturalization. That is, meanings persist not because they have been epistemically earned, but because they have been left unchallenged. A dynamic model of normativity, by contrast, allows for critique, transformation, and the explicit examination of the limits and assumptions of our conceptual schemes. This is where philosophy retains its critical function even within a naturalized framework.

5. Naturalized Normativity as Epistemic Progress

Finally, this enriched version of naturalism allows us to see normativity not as imposed from outside, but as the outcome of epistemic progress. The standards by which we evaluate meaning, knowledge, and interpretation are not static—they emerge from the very practices of inquiry that aim to improve our understanding. Just as Quine emphasized the fallibility and revisability of even our most central beliefs, this model takes normativity to be something constructed, sustained, and revised over time through the lived experience of inquiry.

In this way, meaning stabilizes not merely through adaptive success, but through the self-correcting, historically conscious activity of thinking beings engaged in a shared process of making sense. Without this additional layer of epistemic reflection, our criteria for meaning would be left vulnerable to arbitrariness, ideological reification, or functionalist flattening. But when normativity is understood as the product of an evolving, reflexive, and socially embedded form of inquiry, the gap between fact and norm no longer demands metaphysical closure—it invites philosophical responsibility.

In this context, we may speak of “evolving necessity.” This term, however, requires clarification. It does not aim to reintroduce any metaphysical notion of modality—a move that Quine famously resisted in his critique of the analytic–synthetic distinction and modal logic (QUINE, 1953). Rather, the term refers to a form of epistemic robustness that emerges within historically situated discursive practices. It is closer in spirit to Thomas Kuhn's idea of “normal science,” where certain paradigms become temporarily stabilized through communal agreement and practical success (Kuhn, 1962), or to Robert Brandom's account of materially inferential relations, where normative commitments are articulated and stabilized within a network of reasons and social practices (Brandom, 1994). A cautious reader might worry that speaking of “evolving necessity” risks reinstating a modal vocabulary that Quine explicitly sought to dissolve. However, this account preserves Quine's naturalist continuity by treating necessity not as a metaphysical given, but as a dynamic status that arises from historically sedimented and socially embedded processes of conceptual refinement.

9 Conclusion

At its argumentative climax, this article set out to examine whether the notion of success—whether framed in terms of biological reward structures, behavioral reinforcement, or the stabilization of inferential patterns—can sufficiently explain the transition from mere descriptive regularity to genuine

normativity. In pursuing this inquiry, we have advanced a twofold argument. On the one hand, we acknowledged the strength of Quine's naturalist critique of analyticity and meaning determinacy, which effectively destabilizes any static or metaphysical account of meaning. On the other hand, we explored how teleosemantics attempts to respond to this destabilization by grounding semantic content in the evolutionary success of representational mechanisms.

Yet, our inquiry has shown that success, in its raw biological form, is not enough. Teleosemantics provides a powerful framework for explaining how content *emerges* and *stabilizes*, but not why certain meanings ought to *persist*, be *preferred*, or *guide inference and action* under normative constraints. The article thus answers its central question by arguing that normativity cannot be derived from mere survival or coordination; it must also be shaped by epistemic accountability and reflective practices. It is only through the incorporation of a reflexive, conceptual layer—where systems can assess the *validity* of their representational commitments—that meaning gains its normative traction.

Crucially, this reflexive layer does not require metaphysical postulates or appeals to transcendental realms. Rather, it enables a form of *self-consciousness* that is grounded in our biological background and evolved cognitive architecture. This alignment allows for a stable form of *naturalist empiricism* to persist—one in which normativity emerges not as a rupture from nature, but as a development within it, shaped by our practical engagement with the world and our capacity for conceptual self-correction. By rooting normativity in this biologically plausible reflexivity, we avoid both the arbitrariness of social constructivism and the rigidity of metaphysical apriorism.


The final sections argued that normativity arises where biological function meets epistemic responsibility—not through abrupt shifts or arbitrary cultural constructions, but through a continuous interplay between naturalistic explanation and reflexive self-correction. This interplay, far from being external to science, is an intrinsic part of its evolving structure: a historically situated and socially embedded process capable of producing standards that, while never absolute, aspire beyond local contingencies toward wider epistemic justification. In this light, normativity is not a metaphysical given but an achievement of inquiry—a form of evolving necessity shaped by the demands of understanding, coherence, and critical revision.

By tracing this tension—between evolutionary success and epistemic grounding—we have not only clarified the limitations of purely descriptive accounts but also gestured toward a more sophisticated naturalism. One that, inspired by Quine's continuity thesis, remains loyal to the sciences but insists on philosophy's irreplaceable role in articulating and refining the very standards by which meaning becomes normative. It is in this light that we may also speak of *evolving necessity*: not as a metaphysical constant, but as a form of epistemic robustness forged within historically situated practices of inquiry. Like Kuhn's notion of "normal science" or Brandom's materially inferential norms, necessity here emerges through the iterative processes of stabilization, revision, and communal endorsement that characterize both scientific and conceptual progress. In this light, the article fulfills its aim: not by resolving the problem once and for all, but by showing where and how the solution must begin to take shape.

Rather than presenting a wholesale revision of cognitive science, this perspective aims to complement existing naturalistic models by highlighting how meaning involves more than functional coordination: it involves the capacity to assess, revise, and justify representational commitments over time. What distinguishes meaning in this sense is not just that it works, but that it enables agents to weigh alternatives that may all function well, yet differ in conceptual clarity, inferential coherence, or justificatory strength. Meaning, accordingly, is tied not only to success, but to the capacity for epistemic discrimination within evolving frameworks of inquiry.

References:

- BRANDON, R. *Making it explicit: reasoning, representing, and discursive commitment*. Cambridge: Harvard University Press, 1994.
- CHURCHLAND, Paul M. Eliminative materialism and the propositional attitudes. *The Journal of Philosophy*, v. 78, n. 2, p. 67–90, 1981.
- CHURCHLAND, Patricia S. *Neurophilosophy: Toward a Unified Science of the Mind-Brain*. Cambridge, MA: MIT Press, 1986.
- DAVIDSON, D. Knowing and the cause of action. *Journal of Philosophy*, New York, v. 84, n. 3, p. 131-148, 1987.
- GOODMAN, N. *Ways of worldmaking*. Indianapolis: Hackett, 1978.
- HABERMAS, J. *The theory of communicative action: volume 1: reason and the rationalization of society*. Tradução de T. McCarthy. Boston: Beacon Press, 1984. (Original publicado em 1981).
- KUHN, T. S. *The structure of scientific revolutions*. Chicago: University of Chicago Press, 1962.
- MILLIKAN, R. G. *Language, thought, and other biological categories: new foundations for realism*. Cambridge: MIT Press, 1987.
- MILLIKAN, Ruth G. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press, 1984.
- MILLIKAN, Ruth G. *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press, 1993.
- PAPINEAU, D. The essence of representation: teleosemantics and the selective provenance of meaning. In: *Philosophy of mind and cognitive science*. [S.l.]: Publisher, 2022. p. 14-17.
- QUINE, W. V. O. *Epistemology naturalized*. In: Quine, W. V. O. *Ontological relativity and other essays*. New York: Columbia University Press, 1969. p. 69-90.
- QUINE, W. V. O. *From a logical point of view*. Cambridge: Harvard University Press, 1953.
- QUINE, W. V. O. In praise of observation sentences. In: Quine, W. V. O. *Theories and things*. Cambridge: Harvard University Press, 1990. p. 102-115.
- QUINE, W. V. O. *Ontological relativity and other essays*. New York: Columbia University Press, 1969.
- QUINE, W. V. O. *Progress on two fronts. Proceedings and Addresses of the American Philosophical Association*, Newark, v. 70, n. 2, p. 35-47, 1996.
- QUINE, W. V. O. *Pursuit of truth*. 2. ed. Cambridge: Harvard University Press, 1993.
- QUINE, W. V. O. *The nature of natural knowledge*. In: Gunderson, K. (org.). *Mind and language*. Oxford: Oxford University Press, 1975. p. 67-81.
- QUINE, W. V. O. *The roots of reference*. La Salle: Open Court, 1974.
- QUINE, W. V. O. *Word and object*. Cambridge: MIT Press, 2013.
- QUINE, W. V. O. Pre-established harmony. In: Janssen-Lauret, N.; Kemp, G. (org.). *Quine and his place in history*. London: Palgrave Macmillan, 2016. p. 208-212.
- SMART, J. J. C. Sensations and brain processes. *The Philosophical Review*, Durham, v. 68, n. 2, p. 141-156, 1959. <https://doi.org/10.2307/2182164>.



COGNITIO

Revista de Filosofia
Centro de Estudos de Pragmatismo

São Paulo, v. 26, n. 1, p. 1-15, jan.-dez. 2025
e-ISSN: 2316-5278

 <https://doi.org/10.23925/2316-5278.2025v26i1:e70312>