Teleo-Inferentialism[1]

Ulf Hlobil

Concordia University, Montreal, Canada;

Univerzita Hradec Králové, Czechia

Abstract: The paper presents teleo-inferentialism, which is a novel meta-semantic theory that combines advantages of teleosemantics and normative inferentialism. Like normative inferentialism, teleo-inferentialism holds that contents are individuated by the norms that govern inferences in which they occur. This allows teleo-inferentialism to account for sophisticated concepts. Like teleosemantics, teleo-inferentialism explains conceptual norms in a naturalistically acceptable way by appeal to the broadly biological well-functioning of our innate capacities. As a test-case for teleo-inferentialism, I discuss how the view handles Kripkenstein-style meaning skepticism.

Keywords: meta-semantics; inferentialism; teleosemantics; naturalism; normativity

1 Introduction

The goal of this paper is to sketch a novel meta-semantic theory that combines advantages of normative inferentialism and teleosemantics. I call this theory "teleo-inferentialism." The

---

account conjoins the inferentialist claim that contents are inferential roles with the teleosemantic claim that conceptual norms derive from broadly teleological assessments of capacities as well-functioning or malfunctioning. Unlike traditional normative inferentialism, but like teleosemantics, teleo-inferentialism is naturalistically acceptable. Unlike teleosemantics, but like normative inferentialism, teleo-inferentialism can account for sophisticated contents. The paper is programmatic and my aims are modest: I won't attempt an overall assessment of teleo-inferentialism and how it compares to rival meta-semantic theories. I merely want to sketch the central ideas and show that the view is worthy of further development and investigation.

The paper is structured as follows: In Section 2, I set the stage by explaining some of my assumptions and providing some background on normative inferentialism and teleosemantics, including one central problem for each of these theories. In Section 3, I present my positive proposal, i.e., teleo-inferentialism. And I argue that it solves both problems from the previous section. In Section 4, I show how teleo-inferentialism can handle Kripkenstein-style meaning skepticism. Section 5 discusses two objections, and Section 6 concludes.

2 Background

In this section, I will set the stage. First, I will flag some methodological assumptions. Then I will introduce and compare inferentialism and teleosemantics. For concreteness, I will take Brandom as my stock inferentialist and Millikan as my stock teleosemanticist. However, nothing below hinges on exegetical claims about Brandom or Millikan.

2.1 Methodological Remarks

The teleo-inferentialism I will present is a meta-semantic theory, i.e., it aims to explain what it is in virtue of which contentful items have the contents they have. Sketching any such theory requires that one makes some substantive assumptions. Hence, I want to be upfront about four of my own.

First, I will often talk indiscriminately about linguistic meaning and mental content, assuming that meaningful items in these two domains ultimately have their meanings in virtue of (broadly) the same facts. Teleo-inferentialism is compatible with different claims about the

priority of linguistic over mental content or vice versa. Below, I will often talk as if mental content comes first; this is merely for convenience. I will use "(contentful) item" as a generic term to cover contentful linguistic expressions and the vehicles of mental content.

Second, when I talk about content, I am not interested in the semantic values that we encounter in compositional semantics (see [Yalcin, 2014](#)). Rather, what I mean by "content" is that which allows us to theorize and keep track of the conditions under which what someone believes or asserts is coherent or incoherent, when it follows from something else, when someone has certain (good or bad) reasons for what they believe or assert, and when someone couldn't coherently reject a claim without changing her mind or retracting some of her assertions. In all of these cases, we ascribe content to beliefs, assertions, and rejections; and such ascriptions of content help us to answer questions about the rationality and coherence of the agent's exercises of her conceptual capacities. I will call this kind of content "cognitive role." Perhaps cognitive roles stand in complicated and only loose relations to the information communicated by assertions, truth-conditions, or semantic values in compositional semantics. Perhaps we need other meta-semantic theories for these other kinds of content. I won't address these issues. Rather, I will assume that cognitive role is a good first target for a meta-semantic theory, and I take it to be the traditional target of inferentialism ([Brandom, 1994;](#) [Sellars, 1974,](#) [1953](#)).[2]

Third, in line with normative inferentialism, I will assume that we can give an account of the cognitive roles of meaningful items by specifying their inferential roles.[3] I will thus simply

---

[2] Note that I am not denying that, in the final analysis, cognitive roles should relate in appropriate ways to other kinds of content and to principles of compositionality. I merely leave it to future occasions to spell out such connections.

[3] There are, of course, many kinds of inferentialism (e.g. [Chalmers, 2012;](#) [Block, 1986;](#) [Harman, 1982](#)). Prominent questions about which inferentialists disagree include the following: (i) Are all good inferences meaning-determining or just some privileged subclass? (ii) Are meanings inferential roles, or are they something more traditional (such as extensions or intentions) that is determined by inferential role? (iii) Are the meanings of all meaningful items determined by inferential roles, or are there some meaningful items, perhaps perceptual states, that get their meanings in a different way? (iv) Do inferential roles involve only mental items that are individuated without reference to the environment, or do they involve states of the environment or mental states that are individuated in terms of the environment? (v) Are private mental or overt linguistic uses of meaningful items primary in determining the meanings of these items? The teleo-inferentialism that I am putting forward can stay neutral on all of these points. In effect, for most versions of inferentialism that answer these questions differently, there will be corresponding versions of teleo-inferentialism. But keeping track of the different options goes beyond

assume this and some related core tenets of inferentialism. The only major point over which I disagree with many inferentialists—especially Brandom—will be the nature of conceptual normativity. Since my target is cognitive role, and I take on board the inferentialist claim that we can capture cognitive roles with inferential roles, I will not focus on representational content as teleosemantics does. As far as I can see, teleo-inferentialism is compatible with a teleosemantic account of the representational contents of many items, especially the representational content of conceptually unsophisticated items such as perceptual states. However, this will not be my topic, and I am setting it aside here.

Fourth, following teleosemantics, I assume that it is naturalistically acceptable to appeal to the well- and malfunctioning, or excellence and defect, of capacities of living organisms. That is, I assume that it is naturalistically acceptable to hold, e.g., that it is the function of a heart to pump blood and that, *ceteris paribus*, a heart is malfunctioning or defective if it doesn't pump blood. Of course, it is controversial what biological functions are (see Moosavi, 2019, sec 4). Some think that the function of an organ or property is its causal contribution to a goal or capacity of the overall system of which it is a part (Adams, 1979; Cummins, 1975). Others hold that the function of an organ or property is the causal effect for which it has been selected in the evolutionary history of the organism or a similar selectional history (Godfrey-Smith, 1994; Neander, 1991; Millikan, 1989; Wright, 1973). And still others claim that the function of an organ or property is its usual contribution to the maintenance of a system's organization (Moreno and Mossio, 2015; Mossio et al., 2009; McLaughlin, 2000; Schlosser, 1998). I will stay neutral on such issues and simply assume that we can freely appeal to functions of capacities and, in particular, their well-functioning and defect.

2.2 Two Differences between Inferentialism and Teleosemantics

With these assumptions in place, I want to take a brief look at normative inferentialism and teleosemantics. In particular, I want to bring out two important differences between these theories.

---

the scope of this paper. Hence, I will simply work with Brandom's inferentialism for concreteness without (as far as I can see) relying on any of the features that make it special.

Normative inferentialism, as developed by Brandom (2008; 1994), and teleosemantics, as developed by Millikan (1984), are both heavily influenced by Sellars (1953; 1954; 1967; 1969; 1974). Inferentialism says that possessing a concept is, in the first instance, the ability to engage in rational discourse or thought that involves words or phrases that are governed by certain norms of inference, which Brandom construes as social norms for taking people to be committed and entitled to certain assertions in certain circumstances. Teleosemantics, by contrast, holds that to be a representation is to have the function of playing a certain role in map-like mental structures that are homomorphic to the environment. These characterizations don't even begin to do justice to the two theories, but they allow us to see one commonality and two differences. First the commonality: The two theories agree that content is normative and that this rules out views on which content is fixed by the agent's dispositions to use an item.[4] As Millikan (2005, 82) puts it, "I agree with Brandom that conceptual norms must be disposition-transcendent."

Underneath this superficial agreement, however, there are two important differences. First, Brandom locates the source of conceptual norms in our social practices (Brandom, 1994, xvii, 41, 55, et passim). Millikan, by contrast, thinks that conceptual norms are, at root, of the same kind as norms of well-functioning that govern, e.g., the organs of animals (Millikan, 1984, 17). For Millikan, the relevant conceptual norms are standards of correct representation, about which she writes:

> Enter teleosemantics, with the claim that "representation" is a function term. The evidence? If it were not there would be no standard by reference to which representations would be correct or incorrect, true or false, fulfilled or unfulfilled, satisfied or unsatisfied. [...] That is all there is to the claim that was later dubbed "teleosemantics." (Millikan, 2020, 2-3)

Millikan holds that such standards of correctness, truth, fulfillment, or satisfaction can be explained in terms of the functions of the evaluated representational items. These functions are

---

[4] This claim has come under fire in the last two decades (Glüer and Wikforss, 2015, 2013, 2009b,a; Hattiangadi, 2007, 2006). I have defended normativism about meaning elsewhere (Hlobil, 2015), as have others (e.g. Peregrin, 2012; Whiting, 2009a,b, 2007). For a recent response to my view, see (Reinikainen, 2020). I remain unconvinced by the anti-normativist arguments. But a defense of normativism is beyond the scope of this paper.

fixed by the selectional history of the representations or the system that produces the representations.

Millikan thinks that, in general, the function of representations is to guide behavior or internal processing by providing a map or picture of parts of reality. Put abstractly, Millikan holds that an item, r, represents a part of reality, x, just in case it is used by some interpreting system to guide its activity in a way that requires, for the system's well-functioning, that r plays a role in a larger system of representations, R, that is homomorphic to a larger part of reality, X, and r maps onto x under this homomorphism. For simple representational systems, such as perceptual systems of lower animals, a situation like that is often the direct result of natural selection. An item has been selected for in our evolutionary history, e.g., because the fact that it tracked a part of reality in an appropriate homomorphism yielded effects in the interpreting system that ultimately contributed positively to the overall fitness of the organism. But the situation is more complicated for language and thought, for everyone must acknowledge that conceptual norms are shaped, in some sense, by arbitrary conventions, even if the ultimate source of the norms is biological. To accommodate this, teleosemantics appeals to what Millikan (1984, 18) calls "derived proper functions," to which I will return below.

This brings us to the second difference. In contrast to Brandom, Millikan is primarily interested in representational content and not in what I have called "cognitive role." Indeed, Millikan thinks that the intentionality of beliefs has nothing to do with their cognitive or inferential roles, but rather only with their representational content.

> [T]he intentionality of a belief has nothing whatever to do with its role in inference or indeed with anything that it is supposed to do. [... Rather they] are such because whatever they may ultimately be supposed to do, they cannot do it in accordance with a Normal explanation [i.e., in accordance with their function] unless they map onto something else in accordance with certain rules. (Millikan, 1984, 139)

This passage brings out that Millikan uses conceptual norms in order to explain the mapping between contentful items and the parts of reality that they represent. Brandom, by contrast, uses conceptual norms to explain the cognitive roles of contentful items.

Let's take stock of the two differences between inferentialism and teleosemantics that I have highlighted[5]:

NORM-SOURCE   Inferentialism claims that the normativity of content has its source in our social practices, while teleosemantics claims that this normativity has its source in the (ultimately biological) functions of our organs or capacities (while acknowledging that learning is important).

CONTENT-TYPE   Inferentialism is primarily concerned with content understood as cognitive role, while the kind of content teleosemantics is primarily concerned with is representational, i.e., a mapping of contentful items to parts of reality.

Once we distinguish these two differences, we can see that there is room for at least two neighboring positions. An option that I will ignore is to hold that the normativity of meaning is merely social while targeting primarily representational content. Teleo-inferentialism is a theory of the fourth, remaining kind, i.e., a theory that holds that the normativity of meaning has its source in the—ultimately biological—functions of our capacities while targeting primarily cognitive role. As we will see below, like teleosemantics, teleo-inferentialism acknowledges that social practices play a crucial role in fixing the conceptual norms while also holding that the source of the normativity lies in standards of well-functioning of innate capacities. The four theories can be seen as the four possible choices regarding NORM-SOURCE and CONTENT-TYPE.[6] As we will see below, the choice of teleo-inferentialism is the right choice because it avoids two problems, to which I will now turn.

---

[5] What I say here is similar to Withrington's (2019, 153-154) assessment: "The core feature of teleosemantic theories is the following: they rely on a teleological account of functions to explain the nature of meanings. But most of the available versions of teleosemantics share another feature: they rely on direct links to reality to explain the references and/or truth-conditions of representations." Like Withrington, I will endorse the core feature while rejecting the representational aspect of teleosemantics.

[6] Of course, more options become available if we frame the issue differently; e.g., if we don't insist on the normativity of meaning and content, add other possible sources of this normativity, or add other kinds of content that could serve as our primary target. A discussion of such issues is beyond the scope of this paper.

2.3 Problems of Inferentialism and Teleosemantics

I will argue below that teleo-inferentialism combines advantages of inferentialism and teleosemantics. In particular, like teleosemantics and unlike inferentialism, teleo-inferentialism is naturalistically acceptable and, like inferentialism but unlike teleosemantics, teleo-inferentialism can handle sophisticated concepts. In this subsection, I want to explain these issues for inferentialism and teleosemantics.[7]

Let's start with the insufficient naturalistic credentials of normative inferentialism. Brandom notoriously holds that, in his theory of content, "it is norms all the way down" (Brandom, 1994, 625). And since Brandom doesn't offer any naturalistic account of these norms, many have noticed that it is not clear that naturalists can accept Brandom's inferentialism (Whiting, 2006; Rosen, 2001). DeVries (2013, 260), e.g., writes: "Brandom emphasizes the rational, intentional order but says disappointingly little about its relation to the causal or natural order." And recent advocates of inferentialism often adopt an explicitly non-normative, purely dispositional version of inferentialism because they want to ensure that their meta-semantics is naturalistically respectable (e.g. Warren, 2020b).

Brandom (1994, 28; 2000, 219) thinks, following Kripke (1982), that no such purely dispositional theory can work because dispositions to use contentful items—including dispositions to assess the uses of others—can always include dispositions to make mistakes, even under suitably idealized conditions. But if the conceptual norms that fix contents outstrip (non-normatively specified) dispositions of use, even when idealized (in non-question-begging ways), then naturalist inferentialists need a non-dispositional account of these conceptual norms that is nevertheless naturalistically acceptable. Brandom doesn't have such an account.[8]

---

[7] The two problems that I discuss are, of course, not the only problems for inferentialism and teleosemantics. However, I am setting aside other worries about these theories. For example, I set aside the worry that inferentialism cannot account for the representational dimension of content (Shapiro, 2004), as well as worries about meaning-holism and compositionality (Fodor and Lepore, 1992). Similarly, I am setting aside well-known problems for teleosemantics (Neander and Schulte, 2021, sec 4), such as the worry that teleosemantics cannot explain how there can be distinct representations with (necessarily) coextensive contents (Fodor, 1996), or the worry that Davidson's (1987) Swampman cannot have contentful thoughts. All of these are important worries and must be addressed, but none of them has any special significance for my project of articulating teleo-inferentialism.

[8] Brandom (2008, 12) sometimes appeals to Price's (2004) distinction between subject-naturalism and object-naturalism. The idea is that the question "What is the place of conceptual normativity in nature?" is misleading

Notice that the questionable naturalist credentials of normative inferentialism stem from the questionable naturalist acceptability of the relevant kind of normativity. Teleosemantics doesn't have a similar problem because teleosemantics appeals to a kind of normativity that is in good naturalistic standing (Millikan, 1989).

Let's turn to the problem of sophisticated concepts and capacities for teleosemantics. Neander and Schulte (2021, sec 4.3) write that the "weightiest objection to teleological theories of content" is that "it is unclear how such theories could explain our most sophisticated concepts and cognitive capacities." The problem arises because it is difficult to see how teleosemantics can explain the content of thought and talk, e.g., about democracy, complex numbers, god, or non-actual possibilia. For, if the function of beliefs and assertions about such matters cannot be performed unless these beliefs and assertions are true, the facts that make such beliefs and assertions true or false must make a difference to what we do—and ultimately to the performance of some natural function. But it is far from clear how and why this should be the case, especially for the presumably causally inert representanda above (Peacocke, 1992, 130). Advocates of teleosemantics often acknowledge the problem and leave it for future research. Here is Neander:

> I sympathize with [the] view that we might be unable to explain sophisticated forms of intentionality at present, as its most sophisticated forms might have ties to nonintentional nature that are too complicated and various to be comprehended, at least for now. (Neander, 2017, 6)

While it may be admirable to take one step at a time, this reaction opens the door to the worry that there may not be any naturalistically acceptable account of sophisticated concepts. One approach is to explain sophisticated concepts in terms of simple concepts, e.g., by analyzing sophisticated concepts in terms of simple concepts. Unfortunately, a century of failed attempts

---

because it assumes a representational semantics for the expression "conceptual normativity." If these words don't have any representational job but can be explained entirely in terms of our social practices, then we cannot ask: "Where in nature do we find what we refer to by 'conceptual norms'?" I doubt that Brandom is entitled to this move because he ultimately wants to account for the representational dimension of content. And once he does that, we can pose the question where in nature the referent of "conceptual normativity" can be found. However, I am setting this issue aside here. It suffices for my purposes that, in the eyes of many philosophers, it remains an open challenge to Brandom's inferentialism to show its compatibility with naturalism.

to give necessary and sufficient conditions for sophisticated concepts, like the concept of knowledge, may dampen our optimism regarding such a project. And one may reasonably worry that there aren't many promising alternative approaches to the problem, even in principle.

Notice that teleosemantics's problem with sophisticated concepts stems from its choice to think of content as a matter of representation, i.e., CONTENT-TYPE above. There is no parallel problem for inferentialism. For it is easy to see that sophisticated concepts each have their own inferential role. There are inferences whose correctness clearly turns on the meaning of "democracy" or "god" or concepts for non-actual possibilia.

In this section, we saw that inferentialism's choice of NORM-SOURCE leads to a problem with naturalism and teleosemantics's choice regarding CONTENT-TYPE leads to a problem with sophisticated concepts. The teleo-inferentialism that I will present in the next section avoids these problems by siding with teleosemantics regarding NORM-SOURCE and siding with inferentialism regarding CONTENT-TYPE.

3 Articulating Teleo-Inferentialism

The aim of this section is to articulate teleo-inferentialism. The core idea of teleo-inferentialism is to combine the following two claims:

INF-ROLE   The content of an item is determined by the inferential norms that govern the use of the item.

INF-NORM   The inferential norms that govern the use of an item are ultimately explained by the broadly functional, natural norms that govern exercises of characteristically human capacities.

I take over INF-ROLE from normative inferentialism, and I have just two small things to add.[9] First, if we are interested in contents in the sense of cognitive roles, then INF-ROLE is plausible

---

[9] My commitment to INF-ROLE will be no more and no less problematic than its acceptance by other normative inferentialists, and as already intimated I will here simply assume that this acceptance isn't problematic. A full defense of any normative inferentialist meta-semantics must defend INF-ROLE but I ignore this here because teleo-inferentialism doesn't differ in this respect from other forms of normative inferentialism.

because it cuts contents exactly as finely as we can rationally distinguish them in our inferential practice. This fineness of grain may not be suitable for representational content or informational content or semantic values in compositional semantics, but it is tailor-made for cognitive roles. After all, inferential norms are ideally suited to account for facts about rationality assessments in terms of incoherence, entailment, or possibility of coherently rejecting assertions and beliefs, and these are precisely the facts that we want cognitive roles to explain. Second, INF-ROLE allows us to account for sophisticated concepts. For sophisticated concepts come with sophisticated inferential roles. Thus, teleo-inferentialism shares all advantages of inferentialism over teleosemantics with respect to sophisticated concepts.

This brings us to INF-NORM, which is where most of the action will be. I will use ideas that are broadly similar to teleosemantic appeals to functions in an account of inferential norms. The two crucial moves will be (a) that I will give an account of conceptual norms in terms of inferential capacities and (b) that I will distinguish innate capacities and learned capacities in a way that is broadly similar to Millikan's distinction between primitive functions and derived functions. Let's take these two moves in order.

3.1   Inferential Capacities and Content

In this subsection, I isolate a sense of "capacity" such that capacities come with standards for evaluating their exercises. I then apply this to inferential capacities. I argue that conceptual norms are the standards that come with inferential capacities.

What are capacities? Like dispositions and abilities, capacities are powers, i.e., potentialities inherent in their bearers (Maier, 2010; Bird, 2007). There is a thin and a thick sense of "capacity" in English. In the thin sense "capacity" is roughly synonymous with "potentiality." In this thin sense a bucket may, e.g., have the capacity to hold 10 liters of liquid. By contrast, the thick sense of "capacity" underwrites the inference from "x exercises the capacity to φ" to "x is φ-ing well or badly (or in an entirely failing way), qua exercising the capacity to φ." In this thick sense, any capacity comes with a standard for evaluating its exercises. When I exercise my capacity to ride a bike, say, you can always ask whether this is a poor or a good exercise of that capacity. This distinguishes thick capacities from dispositions that don't supply any standard of evaluation, such as fragility. Henceforth, I will always use "capacity" in this thick sense.

Many abilities are capacities; for many abilities are powers that fix a standard of evaluation. The ability to find your way home, e.g., fixes a standard of success and failure, and even unsuccessful exercises of this capacity can be judged well-performed exercises if the failure is due to external factors that don't count against the exercise, such as unforeseeable construction work that blocks the way. Unlike abilities, however, some capacities can be manifested without any involvement of intentional agency. My immune system has, e.g., the capacity to protect me against antigens. This fixes a standard of evaluation: my immune system can protect me well or badly. But my immune response doesn't involve any intentional agency. Such examples show that the standards of evaluation introduced by capacities do not always depend on anyone's plans or goals, on any ethical norms, or the like.[10] In such cases, the standard that is fixed by a capacity is usually given by the function of the capacity. It is, e.g., the function of my immune system to protect me against antigens, and this is why success in doing this is the standard for evaluating its exercises. As already intimated, I will assume that such talk of functions and the standards for evaluating exercises of capacities is naturalistically respectable, just as advocates of teleosemantics have argued.

Note that capacities come in complex networks. The capacity to write a philosophy paper, e.g., is exercised by exercising many more specific capacities, such as writing paragraphs, quoting correctly, etc. Moreover, the capacity to write a philosophy paper combines with other capacities in non-trivial ways, such as the capacity to come up with interesting ideas and the capacity to formulate them in a compelling way. And when you write a good paper, this may be a way to exercise a general capacity to earn a living. So, capacities and their exercises crisscross and interact in many complex ways. And we should expect this to be true of conceptual and, in particular, inferential capacities.

Let's now turn to inferential capacities, i.e., capacities to infer.[11] The general ability to make good inferences has many species that we can individuate as the inferential capacities to reason

---

[10] Millikan sometimes uses "ability" such that x has the ability to do A iff x is in a state that has the function to result in x doing A (and the state is to some extent intact). She writes: "[W]hat I have an ability to do is what my systems were maintained or selected for doing" (2000, 61). I use "capacity" to express the very similar concept that I explain here because abilities are usually understood as involving intentional agency in a way that neither I nor Millikan (I take it) want to build into the relevant states (Maier, 2010).

[11] I use "reasoning" and "inferring" interchangeably in this paper.

with particular items. So, just as there is the general capacity to dance and the more specific capacity to dance, say, Tango, so there is the general capacity to make inferences and the more specific capacity to make inferences involving, say, "democracy."

Inferring is the exercise of a capacity because it is the manifestation of a potential of the agent that comes with a standard of evaluation. Formally speaking, an exercise of this capacity is good qua exercise of the capacity if it is a good inference (see McHugh and Way, 2018b). Unfortunately, it is controversial what the function of inference or reasoning is. Mercier and Sperber (2017; 2011) hold that the function of reasoning is to enable the agent to engage successfully in argumentative exchanges (see Dogramaci, 2020; Evans, 2011 for criticism). They argue that the hypothesis that the function of reasoning is to produce knowledge or rational beliefs doesn't fit with the overwhelming evidence that reasoning often doesn't produce such results (Kahneman, 2011).[12] McHugh and Way (2018a) argue that the function of reasoning is to yield, ceteris paribus, fitting attitudes towards conclusions, given fitting attitudes towards the premises. And it is often assumed that the function of inference must be closely related to the acquisition of knowledge or true beliefs (Dogramaci, 2017). I suggest that we treat the function of inference as a black-box. We can simply assume that inference—and hence the capacity to infer—has some naturalistically respectable function that provides standards of evaluation.

More specifically, we can assume that each capacity to infer with particular conceptual items also comes with standards for assessing the acts of the capacity. For example, the capacity to reason with "democracy" sets a standards such that reasoning from "This country is a democracy" to "Elections are regularly held in this country" is a good inference, while the inference from "We live in a democracy" to "We live in a society without domination" is a bad act of the capacity, i.e., a bad inference. The standards of these particular capacities are plausibly constraint by the standard of the general capacity for reasoning, and hence the function of reasoning in general. I will say more about how these standards are fixed below. But I first want to bring out the general outlines of teleo-inferentialism.

Let's see where we are: I have assumed above that the kind of content we are interested in, namely cognitive role, is determined by the roles that contentful items play in good inferences.

---

[12] A prominent alternative response is to argue that our reasoning heuristics yield good results in the environments in which they evolved (Gigerenzer and Brighton, 2009).

I have now suggested that inferential capacities have a function that fixes the standards by which their exercises are appropriately evaluated, as the kinds of things they are. Thus, these standards determine which inferences involving a given item are good and which aren't. It follows that these standards fix the content determining conceptual norms and, hence, the content—cognitive role—of the item. So, according to the view that I am putting forward, the conceptual norms are explained by the standards internal to our inferential capacities, which are in turn fixed by the function of these capacities. Since appeal to these functions is naturalistically acceptable, if this account works, it gives inferentialism a naturalistically acceptable notion of conceptual normativity.

Two clarifications are in order. First, one can use an item in deviant ways by exercising capacities that are not the capacity that fixes the content of the item. And perhaps some such uses can be called "reasoning" or "inference." In the overt linguistic case, this might happen, e.g., when someone intentionally presents a bad argument, as an orator might do. Let's call the capacity that fixes the content of an item the capacity that is "characteristically exercised in (inferential) uses of the item." The characteristically exercised capacity is the capacity that one exercises when one reasons with the item in a straightforward and serious way, i.e., without trying to deceive, mislead, be funny, etc. This capacity is what I mean when I speak of "the capacity one exercises when reasoning with an item."[13]

Second, it is worth highlighting that, as already intimated, capacities and their exercises come in complex crisscrossing patterns. Hence, it is not always easy to tell which capacities were used in particular inferences. Take, e.g., the inference from "This fox is a vixen" to "This is a mammal." For one agent, this inference may be an exercise of the capacity to reason from something being a fox to it being a mammal, while for another agent, the inference may be an exercise of the capacity to reason from something being a vixen to it being a mammal. Similarly, someone may reason from "A and not A" to "A" by explosion and someone else by conjunction-elimination. In the first but not in the second case must the agent exercise her capacity to use negation. Perhaps there are cases in which it is indeterminate or vague which

---

[13] The complication that this paragraph addresses is not special to content. When we say that the features that a good hammer must have are determined by the function of a hammer as a hammer, we must focus on the characteristic uses of hammers and exclude uncharacteristic uses of hammers, e.g., those by clowns on stage.

capacities are exercised. In what follows, I will allow myself the idealization that there is always a determinate and unique set of capacities that is characteristically exercised when someone uses a particular item in inferences, i.e., uses it in such a way that the inference's goodness depends on the content of the item.

Before moving on, it is worth pausing to consider two worries. First, an opponent may worry that we cannot understand the function of inference without drawing on a prior understanding of the truth-conditions or reference of the involved items. If we must draw on such a prior understanding, however, then this might threaten our inferentialist ambitions. I take it to be a virtue of teleo-inferentialism that it allows us to clearly identify the need for pure inferentialists to give a naturalistic account of the function of inference that doesn't appeal to ideas such as truth-conditions or reference. And I will return to this issue below. However, I also want to point out that teleo-inferentialism is compatible with a less puristic inferentialism that explains the function of inferences involving simple contents in terms of the representational contents that are supplied by traditional teleosemantics. We could then use teleo-inferentialism to account for sophisticated concepts on that basis. Such a theory would not be a rival to representationalist teleosemantics but a friendly emendation, which helps to account for sophisticated concepts. We can keep an open mind about such options.

Second, an opponent might think that "black-boxing" the function of inference in a justification of INF-NORM has all the characteristic "advantages of theft over honest toil" (Russell, 1919, 71). I have two things to say in response. Firstly, every naturalist should hold that there is a correct naturalistic explanation of the norms of good inference. For, the naturalist cannot appeal to inferential norms that don't have such an explanation, on pain of giving up her naturalism. And rejecting that there are any inferential norms isn't a stable position, at least not for anyone as heavily involved in inferential practices as a philosopher. Secondly, there is some distance between the claim that inference has a general function and the claim that inferential norms governing particular concepts can be explained in terms of broadly biological norms of well-functioning and defect. And the honest toil I have to offer in the next subsection will go into traversing that distance.

3.2 Learning Inferential Capacities

In the previous subsection, I offered a very abstract picture of naturalistically acceptable inferential norms, which I suggested we can use in an inferentialist meta-semantics. The idea was that our conceptual capacities fix the inferential norms and thereby the contents of the items with which we reason. However, our capacities for inferring with particular, perhaps sophisticated, concepts are not innate and cannot be fixed by our genetic makeup or the like.[14] The capacities to reason correctly regarding democracy or complex numbers, e.g., are acquired by learning. And the conceptual norms governing "democracy" and "complex number" are obviously a matter of arbitrary convention, at least in the sense that we could have different signs to carry the contents that "democracy" and "complex number" actually carry. So how can we understand the norms internal to these capacities on the model of norms fixed by biological functions? The aim of this subsection is to answer that question.

The basic idea behind my answer is this: Our inferential capacities are the results of the exercise of a more general and innate capacity, namely the capacity to learn inferential capacities. And I suggest that the standards of particular inferential capacities can be explained in terms of the standards that are internal to the capacity to learn inferential capacities.

How we should think about the capacity to learn inferential capacities will depend on what exactly we take inference to be and how we can acquire intellectual capacities. It will depend, e.g., on whether we think of inferential capacities as primarily capacities for mental acts or for public linguistic acts, and it will depend on what roles direct experience with objects have relative to linguistic exchanges with caregivers, etc. Since most of these issues strike me as ultimately

---

[14] A referee points out that such talk could be misunderstood as suggesting that we can have a particular concept before we can reason in a way that is governed by the norms for that concept. That would indeed be a misunderstanding. My view is that we reason in ways that are governed by certain norms, and that we have a particular concept if and because we engage in reasoning that is governed by the appropriate norms. Nevertheless, I hold that we reason with concepts. Compare (that old warhorse) chess (which is still a source of great examples, and less parochial than games like baseball): We play with chess pieces, even though they are chess pieces only in virtue of being governed by certain norms. While there is a sense in which it is true that we play with pieces of wood, we don't play with pieces of wood as pieces of wood but rather with pieces of wood as chess pieces. Similarly, while it is true that we reason with the vehicles of contents such as inscriptions or sounds or the like, we don't reason with them as inscriptions or sounds or the like but as concepts.

empirical questions, I think that our philosophical meta-semantics should, as much as possible, stay neutral with respect to them.

Fortunately, it should be uncontroversial that we have a capacity to learn inferential capacities. After all children usually acquire capacities to reason with particular contentful items (however imperfectly). Hence, they must have had the potentiality to learn these particular inferential capacities. Moreover, this potentiality to learn inferential capacities can be manifested well or poorly. When one forms, e.g., the disposition to infer "These fact don't support claim *C*" from "These facts support claim not-*C*" (without one's learning opportunities suggesting that this is a good inference), this is a defective exercise of one's capacity to learn the capacities to reason with the item "support." So the potentiality to learn inferential capacities is a capacity. It comes with its own standards of well-functioning. This suffices for my purposes; I can remain neutral on the difficult (ultimately empirical) issues mentioned above.

How does the capacity to learn inferential capacities explain the standards internal to particular inferential capacities? I submit that the capacity one characteristically exercises when one reasons with an item is the capacity that would result from a non-defective exercise of the capacity to learn inferential capacities, given the learning opportunities one had. For example, when I am reasoning with "+", I am exercising my capacity for addition.[15] And this is so because, given the learning opportunities I had in the past, a non-defective exercise of my capacity to learn inferential capacities would result in a non-defective capacity to reason with "+" in accordance with the norms governing addition.

Given our finite and fallible nature, we hardly ever exercise any of our capacities in an entirely non-defective way. This applies to our capacity to learn inferential capacities. Put differently, even under ideal learning conditions, we hardly ever acquire an entirely non-defective capacity to infer in certain ways. For the addition example, this means that we usually don't acquire an entirely non-defective capacity to reason with "+" in accordance with the norms governing addition. Nevertheless, the capacity that we acquire is a capacity for addition,

---

[15] I will assume throughout the paper that we in fact mean addition by "+". As will become clear in the section on Kripkenstein below, on my view, this is a substantive claim about what would be the result of an entirely non-defective exercise of the characteristically human capacity to learn inferential capacities, given our typical learning opportunities with respect to "+". I will assume that such an entirely non-defective exercise of our innate learning capacity yields a capacity to add numbers.

albeit an imperfect instance of this capacity. This explains why the exercises of this capacity are governed by the norms for addition, even if we are disposed to make mistakes.

At this point, it may be helpful to clarify these ideas further by offering a comparison with similar ideas familiar from teleosemantics. Millikan ([1984;](#) [1993](#)) wants to allow broadly biological norms to apply to exercises of learned capacities. And like the account I am suggesting here, she wants to explain the norm-governed-ness of the learned capacities in terms of the innate capacities.

> In the case of innate abilities, no matter what dispositions a mechanism happens to have, what determines its abilities is what it was selected for doing. In the case of learned abilities, what natural selection selected for was the ability to learn in a certain way. It selected for mechanisms that became tuned through interaction with the environment to do things of useful kinds. For an organism to know how to do *A* as a result of learning is for it to possess an intact mechanism that is biologically designed to be tuned to do things like *A* and that has been tuned to do *A* as designed. That is, it became tuned in the same manner, following the same principles, as its successful ancestors when they were learning to do similar kinds of things. ([Millikan,](#) [2000,](#) 63)

At this level of abstraction, my account is very similar to Millikan's. On my view, our inferential capacities are tuned, by learning (which is an interaction with the environment), in such a way that we can reason with given items in particular ways (which is a thing of a useful kind). And for us to know how to reason with, say, "+" as a result of learning is for us to possess an intact capacity that is (broadly) biologically designed to be tuned to produce capacities like the capacity to reason with "+" and that has been tuned to reason with "+" as designed. That is, our general and innate capacity to learn reasoning capacities became tuned (i.e. produced a particular inferential capacity) in the same manner, following the same principles, as its successful ancestors in the acquisition of similar inferential capacities.

There are also, however, two important differences between Millikan's view and mine. First, given my focus on cognitive role, rather than representation, I don't think that the relevant biological function is to guide behavior or cognitive subsystems by mapping parts of reality. Rather, I think that the relevant general function is that of inferring well. As explained above, I want to stay neutral regarding the function of inference. But whatever the function of inference

is, our ability to acquire inferential capacities has the function of allowing us to acquire particular capacities that fulfill the function of inference.

The second difference has to do with Millikan's idea of derived functions, to which she appeals in a footnote to the just-quoted passage (Millikan, 2000, 63fn3). Millikan holds that learned abilities have proper derived functions. She says that an item *A* has function *F* as a "proper derived function" iff item *A* originated "as the product of some prior device that, given its circumstances, had performance of *F* as a proper function and that, under those circumstances, normally causes *F* to be performed by means of producing an item like *A*" (Millikan, 1993, 13-14). If that were my account, the function of the capacities to learn inferential capacities and the function of these learned inferential capacities would have to be identical. But that seems false. The function of the capacity to learn inferential capacities is to produce non-defective inferential capacities. The function of inferential capacities, by contrast, is to produce good inferences. This is important because if there weren't two distinct kinds of function here, we would have only one normative standard. In fact, however, the capacity to learn inferential capacities comes with its standard of defect and well-functioning; and particular inferential capacities have each their own standards of defect and well-functioning. That is what allows us to say that there are inferential norms, and hence cognitive roles, for different items that are distinct from each other and not fixed by our innate biological makeup.

With this comparison to Millikan's view, I hope to have clarified the teleo-inferentialist account of inferential capacities in two respects: First, like Millikan, I hold that the conceptual norms that determine content are fixed by the result of non-defective exercises of our capacity to learn concepts, which for me just are non-defective exercises of our capacity to learn inferential capacities. The relevant conceptual norms are fixed by the capacities that result from non-defective exercises of our learning capacities, given the learning opportunities with which we are presented.

Second, and in contrast to Millikan, I hold that the learning opportunities don't just tune an innate device to perform its biologically fixed function. Rather, I hold that genuinely new particular functions can arise. All these new functions are particular instances of the general function of inferring well. But this general function divides into many genuinely distinct functions for distinct concepts, i.e., distinct items with which we can reason.

To clarify my ideas further, consider the following analogy: We may suppose that we have an innate and biologically fixed capacity to learn social norms,[16] such as norms of politeness or morality. Indeed, anthropological research suggests that this capacity is one of the most striking features of our human nature (Tomasello, 2019; Schmidt et al., 2019, 2012). The function of this general capacity, we may suppose, is to allow us to live together in a well-coordinated way. The capacity to follow a particular norm of politeness or morality, however, has the function of yielding acts that are in accordance with this norm of politeness or morality. In this way, the general function of our innate capacity to learn social norms gives rise to particular functions of particular capacities, and those correspond to particular norms. Which norms those are depends, of course, on the social norms that we learn. In this sense, it is our learning environment and not our biology that determines the norms. But the environment can only do this because of our innate and broadly biological capacity to learn such norms, and this innate capacity is governed by its own internal norms. My claim is that the case of inferential capacities is of this kind. We learn particular inferential norms, and those are determined by the inferential norms in use in our environment. But we can only do this because we have a general and innate capacity to learn inferential capacities.

This concludes my exposition of teleo-inferentialism. To sum up, according to teleo-inferentialism, the content, i.e. cognitive role, of an item is determined by the inferential norms that govern its use. These norms are, in turn, fixed by the capacity that is characteristically exercised in reasoning with said item. The capacity we characteristically exercise when reasoning with a given item is the capacity such that a non-defective instance of it would be the result of a non-defective exercise of our innate and (broadly) biological capacity to learn inferential capacities, given our learning opportunities. The standards of well-functioning for this general learning capacity are correlated with the function of inference, but we can black-box that function and leave it for future research. Teleo-inferentialism gives an account of cognitive roles in terms of inferential roles. And sophisticated concepts can be characterized by their sophisticated inferential roles. Hence, unlike teleosemantics, teleo-inferentialism doesn't run

---

[16] If we use a wide notion of social norm, on which inferential norms count as social norms, then this isn't really an analogy but rather the general case of which the case of inferential norms is merely a particular instance. Thanks to a referee for prompting me to clarify this.

into any problems with sophisticated concepts. Unlike Brandomian inferentialism, teleo-inferentialism is naturalistically respectable: while it leaves considerable room for the importance of social practices, it locates the ultimate source of conceptual norms in the (broadly) biological function of our capacity to learn inferential capacities, which in turn depends on the function of inference.

3.3   Worry: Can't We Learn the Right Thing in the Wrong Way?

Before moving on, I want to address an objection.[17] On the view I just presented, you cannot acquire a capacity to reason in a certain way by exercising your capacity to learn inferential capacities in a defective and deviant way. And in general, I claim that you cannot acquire the capacity to do X by defectively exercising your learning capacities in circumstances where a non-defective exercise of these learning capacities would yield a capacity to do Y.

But it can seem that there are counterexamples to such a principle. Suppose, e.g., that in a Robinson Crusoe style situation, a parent teaches their child, Abby, to spell English words in the British way, but the child exercises her learning capacity in a defective way, and as a fluke occurrence, this defective way of exercising her learning capacity yields a robust disposition to spell English words perfectly in accordance with the American spelling.[18] Perhaps Abby is later rescued, after her parents have died, and ends up in America without ever encountering any negative evaluations of her English spelling. My theory of the norms governing learned capacities, commits me to saying that Abby has a defective capacity to spell the British way and not a non-defective capacity to spell the American way, when she arrives in America. Isn't that an obviously false implication of my view?

I endorse this consequence of my view. The temptation to say that Abby has and exercises a capacity to spell the American way is that she conforms perfectly to the norms of American spelling. But merely conforming to a norm isn't sufficient for exercising a capacity whose internal norm matches the norm with which one conforms. After all, there are always infinitely many abstract norms to which one conforms perfectly. So if Abby really exercises a capacity to

---

spell in the American way, this must be because her new social context in America makes it the case that the norms for spelling in the American way are now the ones relevant for determining which capacities she has and exercises.

I hold that Abby will acquire the capacity to spell in the American way if she lives in America for a while and that this acquisition of a new capacity may be unnoticeable to Abby herself. But this is because we can repurpose aspects of our capacities in learning new capacities, and we need not be aware of this for it to happen. When Abby exercises her imperfect capacity to spell the British way in her new environment, she will not encounter any negative consequences and this will implicitly encourage her to continue spelling words in the way she does. Once this happened for a long enough time, I think we can describe Abby as having acquired a new capacity by repurposing her old capacity without ever noticing that this was going on. In Millikan's terminology, Abby's "spelling mechanism" is now "tuned" to American spelling, it is merely a limiting case as no actual adjustments were necessary (like when you tune a guitar and realize that it is already tuned).

4   A Teleo-Inferentialist Response to Kripkenstein

In the previous section, I have presented teleo-inferentialism, and I have argued that teleo-inferentialism is naturalistically acceptable while also applying to sophisticated concepts. As already intimated, a full comparison and defense of teleo-inferentialism with respect to alternative meta-semantic theories is beyond the scope of this paper. It is still useful, however, to have a look at how teleo-inferentialism performs with respect to an infamous meta-semantic problem, namely Kripkenstein's meaning-skepticism (Kripke, 1982). That is the aim of this section.

Kripke's Wittgenstein (aka Kripkenstein) famously asked in virtue of what fact I meant in the past addition by "+" rather than quaddition, where quaddition agrees with all uses of the sign "+" that I ever encountered. For concreteness, we can assume that I only encountered additions of numbers smaller than 56 and that quaddition (written "$\oplus$") is defined as follows:

$$x \oplus y = \begin{cases} x + y, \text{ if } x, y < 57 \\ \quad 5 \text{ otherwise} \end{cases}$$

22

Dispositionalist accounts say that I meant addition in virtue of having a certain disposition (Warren, 2020a). Kripke (1982) argues that dispositional accounts fail for three reasons: (i) Dispositions are finite but our meaning must determine infinitely many cases. (ii) We have dispositions to make mistakes. (iii) Meaning is normative but dispositions are not.

Teleo-inferentialism handles Kripkensteinian skepticism as follows: The fact in virtue of which I meant addition rather than quaddition by "+" is that my uses of "+" were governed by the inferential norms that determine the meaning of addition, which we may take to include, e.g., some inferential-rule versions of the Peano axioms (see Warren, 2020b, 200). My uses of "+" were governed by these norms because they were exercises of my capacity to reason with "+," and this capacity determines the standards that apply to its exercises, which in my case were inferential norms that fix the addition function, such as the inferential-rule versions of Peano Arithmetic.

Why was my capacity a capacity to perform additions rather than a capacity to perform quaddition? It was a capacity to perform additions because my learning opportunities triggered in me a capacity to learn addition, rather than a capacity to learn quaddition. That is, a non-defective exercise of my capacity to learn inferential capacities as applied to my learning opportunities regarding reasoning with "+" would yield a non-defective capacity to reason with "+" in accordance with the inferential norms for addition. In other words, it is part of my human nature that the capacity that I learn by the kind of training I received is a capacity for addition and not a capacity for quaddition. This is because a well-functioning human being would engage in non-defective exercises of her capacity to learn inferential capacities and, given human nature, this would result in a capacity for addition.[19] That is simply a natural fact about humans.

---

[19] The account I am offering here has similarities to Millikan's response to Kripkenstein, where she writes: "[U]nless doing arithmetic results from a total breakdown of the cognitive systems [...] then whatever you mean to do when you encounter 'plus', that content has been determined by your experiences coupled with evolutionary design. But reasonably, whatever you mean by 'plus' is the same as what other people mean who are endowed with the same general sort of cognitive equipment and have been exposed to the same sort of training in arithmetic. This meaning has been determined by the application of Homo sapiens rules [i.e., rules whose following is an innate function (I would say capacity) in humans] of some kind to experience. It is likely that these are extremely abstract general purpose Homo sapiens rules, in accordance with which human concept formation takes place" (Millikan, 1990, 343). The main difference is that I am more relaxed about tying innate human capacities to the environment and evolutionary history. This is mostly a consequence of my focus on cognitive roles, rather than representations.

It is not that we could not learn quaddition. But in humans with a non-defective capacity to learn such things, the training to learn quaddition would have to look very different from the training we receive when learning addition. It would have to mention, e.g., that cases in which one of the numbers is larger than 56 are special. Special in what way? Special in the way that they must be given a special treatment in training that leads to a capacity for quaddition. It is special in being not the kind of thing we humans take for granted in learning capacities, at least not if our capacities to learn capacities are well-functioning.[20] There is nothing special in that way about the capacity to add, relative to our typical learning opportunities for reasoning "+." That is why we learn addition.[21]

A human with defective capacities to learn capacities may acquire a disposition to reason in accordance with quaddition upon receiving the usual training in addition, perhaps with this training being restricted to numbers less than 57. Such a person would thereby acquire a defective capacity to use "+" and, hence, a defective capacity for addition. In a Martian in whom such a (we may even assume physically indistinguishable) capacity to learn capacities is non-defective, the resulting disposition would count as a capacity for quaddition. However, in a human the capacity is not a capacity for quaddition but rather a defective capacity for addition. This brings out that whether the resulting disposition is a capacity for addition or quaddition does not depend on the history that brought it about or the physical constitution that underlies it. Rather, it depends on the nature of the bearer. If the bearer is a human, then certain things

---

[20] Although I don't want to argue this exegetical point here, this is arguably a position that can be found in Wittgenstein. Consider, e.g., the following passages (translations are mine): "Does it fit our needs to count as we count? That's how we live. It belongs to human natural history" (Wittgenstein, 2000, item 137, 60b-61a). "But is the function of negation in reports and rejections of orders ('Don't do that!') the same?—What we call, or don't call, the same function will depend on human nature. Just like what is, or isn't, a basic need. [...] Human nature determines what is capricious" (Wittgenstein, 2000, item 137, 111a). "It lies in human nature to understand finger-pointing like that [i.e. in the direction of the finger-tip]" (Wittgenstein, 2000, item 155. 82r).

[21] On this point, I agree with Lewis's (1983, 376) response to Kripkenstein: "How ironic that we were worried to find nothing positive to settle the matter in favour of addition! For the lack of anything positive that points either way just is what it takes to favour addition. Quaddition, being less natural and eligible, needs something positive in its favour. Addition can win by default." The difference is that, unlike Lewis, I don't hold that addition is in itself more natural than quaddition. I need only the weaker claim that the capacity to add is the one that is naturally acquired by the training we typically receive in addition, i.e., it lies in our human nature to acquire this capacity if we exercise our capacity to acquire capacities in a non-defective way upon encountering the typical training.

count as well-functioning and others count as defective. And similarly for the Martian. And that is what makes the disposition a capacity for addition in the human and a capacity for quaddition in the Martian.

This teleo-inferentialist response to Kripkenstein can acknowledge the three facts that Kripke uses to rule out dispositional theories. First, teleo-inferentialism can acknowledge that my dispositions to use "+" cover only finitely many cases while the meaning of "+" covers infinitely many cases. It is true in general that a capacity to do something fixes a standard of evaluation for infinitely many cases, and this often includes far-out cases. The capacity to read a book aloud, e.g., sets a standard of evaluation that applies to arbitrarily long books, even though no human being could read a book over a certain length.

Second, teleo-inferentialism can acknowledge that I have dispositions to make mistakes in additions. In general, capacities can be possessed imperfectly. I have dispositions to make (even systematic) mistakes when reading aloud. That doesn't show that I don't have the capacity to read aloud, or that I exercise a different capacity when I read aloud.

Third, teleo-inferentialism explicitly acknowledges the normativity of meaning. My understanding of addition guides me in new applications of "+" because this understanding consists in my capacity to reason with "+" and what I do is guided by that capacity. Note that this is not the guidance provided by a premise in practical reasoning, as some foes of conceptual normativity have assumed it must be (Glüer and Wikforss, 2009a, 2015, 2013). Rather our uses of concepts are guided by our capacities in the sense in which the raising of my arm is guided by my capacity to raise my arm. That suffices to ensure that my use of "+" in a new case is not an "unjustified leap in the dark" (Kripke, 1982, 10).

An opponent may object that, for all we know, given my learning opportunities, a non-defective exercise of my capacity to learn inferential capacities may have resulted in a capacity for quaddition and not addition. Hence, my current use of "+" may be massively mistaken because I really always meant quaddition by "+." The response is that there is a fact of the matter which capacity would result from a non-defective exercise of the capacity to learn inferential capacities. I may be mistaken about which one that is, but that there is a fact of the matter suffices to answer the meaning skeptic.

Of course, it can also happen that the learning opportunities do not suffice to uniquely determine the capacity that would result from a non-defective exercise of the capacity to learn

inferential capacities. Perhaps we have too few examples. Such cases will lead to situations in which two people with the same learning opportunities can rationally mean different things by the same item or have incomplete grasps of meanings. The crucial point is, however, that the learning opportunities we normally encounter for addition are not insufficient in this way. In this case, the learning opportunities together with human nature fix the meaning uniquely to be addition (or so we assume when we worry about Kripkenstein).

The opponent might reply that even a non-defective exercise of the capacity to learn inferential capacities can result only in a finite inferential capacity, which is compatible with many infinitary inferential roles. That worry misses that capacities come with normative standards and that these normative standards are not restricted to finitely many cases. Teleo-inferentialism holds that inferential capacities are naturalistically intelligible states of animals and also that these capacities fix standards of evaluation for all possible exercises of that capacity. As theorists we can capture these infinitary norms by finite means by specifying them recursively or schematically, as we do when specifying inferential rules in logic. For the case of addition, we can specify the norms by giving inferential rules that capture the Peano Axioms of arithmetic (Warren, 2020b, 200). There is nothing mysterious about this feature of inferential capacities. My capacities to drive a car, to cook omelets, or to speak grammatically, and the like all apply to infinitely many cases and determine standards of evaluation for all these cases.

The opponent may reply that there really is no such fact of the matter because no naturalistically acceptable description of my capacity to learn inferential capacities suffices to ensure that, given my learning opportunities, any exercise of the learning capacity that didn't yield a capacity for addition is defective. While that may be true if we are not allowed to describe the capacity to learn inferential capacities in normative terms, it is obviously false if we can so describe it. But I have assumed that we can freely appeal to evaluative notions, such as "being non-defective only if it yields a capacity for addition." According to the kind of naturalism that I wish to defend, capacities are parts of nature and this acknowledgment isn't hostage to the possibility of reducing capacities to more basic parts of nature.[22]

---

[22] Here I am signaling my openness to what is sometimes called "liberal naturalism," i.e., the view that nature may include more than is assumed by the natural sciences (Beasley, 2020; Macarthur, 2014, 2019, 2004). Of course, whether I am forced into such a view will depend on whether natural science can accept unreduced talk of

I conclude that teleo-inferentialism can handle Kripkensteinian skepticism about meaning. This shouldn't be surprising. Kripke already suggested that the core of the issue is the normativity of meaning, and teleo-inferentialism explicitly acknowledges this normativity. What should be good news for naturalists, however, is that teleo-inferentialism allows us to acknowledge the normativity of meaning and, hence, solve Kripkensteinian worries in a naturalistically acceptable way.

5 Two Objections

In this section, I want to discuss two potential objections to teleo-inferentialism. The first objection says that teleo-inferentialism cannot accommodate incomplete understanding of contents. The second objection says that inference doesn't have any natural function.

Although I have above set aside general worries about inferentialism, the first objection is a worry about conceptual role theories (and some other theories) in general. I want to discuss this worry because I think it is particularly instructive to see how teleo-inferentialism offers a response. Greenberg has recently argued that conceptual-role theories, such as inferentialism, and covariation theories, such as teleosemantics, have a problem with the fact that we can have an incomplete understanding of contents, i.e., a thinker can have "a concept without having beliefs or inferential dispositions sufficient to individuate the concept" ([Greenberg, 2014,](#) 150). With respect to inferentialism, the problem is supposed to be that inferentialism cannot distinguish two cases in which agents have the same inferential dispositions regarding a particular item but these dispositions constitute an incomplete understanding of one concept, in the first case, and a perfect understanding of another concept, in the second case. For, if the content grasped is determined by the inferential dispositions, then the item must have the same content in both cases.

> [A]ccording to conceptual-role theories, to have a concept is to have the concept's canonical disposition. But, assuming the phenomenon of incomplete understanding is genuine, one

---

capacities, and whether I really need to appeal to capacities as unreducible parts of nature. Addressing these issues is beyond the scope of this paper. This issue is why I regularly qualify "biological" by "broadly" in the text.

who has a concept need not be disposed to make any particular inference. (Greenberg, 2014, 151)

Teleo-inferentialism is not subject to this problem. For, according to teleo-inferentialism, what determines the content of an item are not the agent's inferential dispositions but the capacity that the agent characteristically exercises when reasoning with the item. Suppose, e.g., that someone has flawlessly exercised her capacity to learn how to reason with an item, *a*, in accordance with inferential role *A*, and someone else has poorly exercised his capacity to learn how to reason with an item, *b*, in accordance with inferential role *B*. Suppose also that the dispositions of our two agents to use items *a* and *b*, respectively, are identical, including their categorical bases. According to teleo-inferentialism, items *a* and *b* can still differ in their content, as long as *A* and *B* are distinct. Our first agent has a complete understanding of content *A*, while our second agent has an incomplete understanding of content *B*. Hence, teleo-inferentialism can accommodate the phenomenon of incomplete understanding.

A second objection to teleo-inferentialism may be that inference doesn't have any natural function. If that is right, then it is implausible that we have a capacity to learn inferential capacities that comes with broadly biological norms of well-functioning and defect. Dutilh Novaes (2018, 516) suggests this much when she urges us to take seriously the "hypothesis that reasoning may be a product of nonselective forces in evolution (e.g., an exaptation), or perhaps not a product of biological evolution at all."[23] My response is that teleo-inferentialism does not imply that the norms that are internal to our capacity to learn inferential capacities and, in turn, to those inferential capacities must be explained in terms of adaptation. Teleo-inferentialism merely needs the two weaker claims that (i) there are such norms and that (ii) they can be understood in a naturalistically acceptable way. That there are such norms is plausible because we don't hesitate to evaluate these capacities as well-functioning or defective. We say that people

---

[23] Dutilh Novaes (2021) also argues in Chapter 10 of her recent book that deduction is not an adaptation. However, she also holds that deduction is a very special kind of practice and special kind of reasoning. Hence, the relevance of her claim about deduction for the question whether inference or reasoning in general are adaptations is not immediately clear. So I focus on Dutilh Novaes's claim about reasoning in general. Dutilh Novaes's claims are strongly influenced by Heyes's (2018) work on cognitive gadgets. So it may be worth mentioning that Heyes is happy to say that evolution has provided us with general-purpose mechanisms of learning and memory. It may well be that the capacity to learn inferential capacities is a general-purpose mechanism of a similar kind and generality. Thus, teleo-inferentialism may be compatible with many of Dutilh Novaes's and Heyes's views.

have learning disabilities if they have severe problems with learning inferential abilities. And in less severe cases, we may call them "slow learners." These are clearly ascriptions of defects. Similarly, we think that there is a better and worse grasp of concepts that one may have, and a poor grasp of a concept is a defective inferential ability, according to teleo-inferentialism. Hence, we engage in many evaluations that presuppose that the norms posited by (i) exist.

As for the claim (ii) that these norms are naturalistically intelligible, this claim is, of course, non-optional for a naturalist who accepts (i). And I take my target audience here to be already committed to naturalism, at least in a broad sense.

More generally, it seems to me that inferential capacities and the grasp of concepts are an essential part of a good human life. So it is also essential to a good human life to learn such inferential capacities. This standard of a good human life will suffice for teleo-inferentialism. Naturalists will want to give a naturalistically acceptable account of what a good human life is. Perhaps the adaptation of our capacities to learn how to reason won't play any part in that account. That won't undermine teleo-inferentialism as long as a good human life requires that we learn inferential capacities, and hence can do so well or poorly.

6   Conclusion

The aim of this paper was to put teleo-inferentialism on the table. Teleo-inferentialism is a meta-semantic theory that combines the naturalism of teleosemantics with the inferential-role semantics of normative inferentialism. It says that contentful items have their contents in virtue of the norms that are fixed by the capacities that are characteristically exercised in uses of these items in reasoning. What capacities these are depends, in turn, on the more basic capacity exercised in acquiring the more complex capacities. This bottoms out in an innate, natural capacity to learn inferential capacities.

Many questions remain open: How do the contents ascribed by teleo-inferentialism relate to the semantic values in compositional semantics? Where does teleo-inferentialism stand on standard worries about inferentialism, such as worries about meaning-holism, compositionality, and semantic externalism? Can teleo-inferentialism account for the representational dimension of content? Etc. All such questions will have to wait for future occasions.

I hope, however, to have shown that teleo-inferentialism is a novel option on the menu of meta-semantic theories and that it has some distinctive advantages. In particular, teleo-inferentialism can account for inferential norms in a naturalistically acceptable way. And in doing so, teleo-inferentialism isn't limited to contents that have close connections to perceptual contents. Rather, teleo-inferentialism can draw on all the complexity of our social, discursive practices, as we encounter them when we learn new concepts, in order to explain sophisticated contents.

References

Adams, F. R. (1979). A goal-state theory of function attributions. *Canadian Journal of Philosophy*, 9(3): 493–518.

Beasley, B. (2020). Naturalism without a subject: Huw Price's pragmatism. *Inquiry*, doi: 10.1080/0020174X.2020.1820903

Bird, A. (2007). *Nature's Metaphysics*. Oxford University Press, Oxford.

Block, N. (1986). Advertisement for a semantics for psychology. *Midwest Studies in Philosophy*, 10(1): 615–678.

Brandom, R. B. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Harvard University Press, Cambridge, Mass.

Brandom, R. B. (2000). *Articulating Reasons: An introduction to inferentialism*. Harvard University Press, Cambridge, Mass.

Brandom, R. B. (2008). *Between Saying and Doing: Towards an analytic pragmatism*. Oxford University Press, Oxford.

Chalmers, D. (2012). *Constructing the World*. Oxford University Press, Oxford.

Cummins, R. (1975). Functional analysis. *Journal of Philosophy*, 72: 741–64.

Davidson, D. (1987). Knowing one's own mind. *Proceedings and Addresses of the American Philosophical Association*, 60(3): 441–458.

Devries, W. A. (2013). All in the family. In Ryder, D., Kingsbury, J., and Williford, K., editors, *Millikan and Her Critics*, pages 259–280. Wiley.

Dogramaci, S. (2017). Why is a valid inference a good inference? *Philosophy and Phenomenological Research*, 94(1): 61–96.

Dogramaci, S. (2020). What is the function of reasoning? On Mercier and Sperber's argumentative and justificatory theories. *Episteme*, 17(3): 316–330.

Dutilh Novaes, C. (2018). The enduring enigma of reason. *Mind and Language*, 33(5): 513–524.

Evans, J. S. B. (2011). Reasoning is for thinking, not just for arguing. *Behavioral and Brain Sciences*, 34(2): 77–78.

Fodor, J. A. (1996). Deconstructing Dennett's Darwin. *Mind and Language*, 11(3): 246–262.

Fodor, J. A. and Lepore, E. (1992). *Holism: A Shopper's Guide*. Blackwell.

Gigerenzer, G. and Brighton, H. (2009). Homo Heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, 1(1): 107–143.

Glüer, K. and Wikforss, A. (2009a). Against content normativity. *Mind*, 118(469): 31–70.

Glüer, K. and Wikforss, A. (2009b). Against normativity again: reply to Whiting. Unpublished manuscript. Accessed October 10th, 2009. URL: http://people.su.se/%7Ekgl/Reply%20to%20Whiting.pdf

Glüer, K. and Wikforss, A. (2013). Against belief normativity. In Chan, T., editor, *The Aim of Belief*, pages 80–99. Oxford University Press, Oxford.

Glüer, K. and Wikforss, A. (2015). Meaning normativism: Against the simple argument. *Organon F*, 22(1): 63–73.

Godfrey-Smith, P. (1994). A modern history theory of functions. *Noûs*, 28(3): 344–362.

Greenberg, M. (2014). Troubles for content. In Burgess, A. and Sherman, B., editors, *Metasemantics: New Essays on the Foundations of Meaning*, pages 147–168. Oxford University Press, Oxford.

Harman, G. (1982). Conceptual role semantics. *Notre Dame Journal of Formal Logic*, 28: 242–56.

Hattiangadi, A. (2006). Is meaning normative? *Mind & Language*, 21(2): 220–240.

Hattiangadi, A. (2007). *Oughts and Thoughts: Rule-Following and the Normativity of Content*. Oxford University Press, Oxford & New York.

Heyes, C. (2018). *Cognitive Gadgets*. Harvard University Press, Cambridge, Mass.

Hlobil, U. (2015). Anti-normativism evaluated. *International Journal of Philosophical Studies*, 23(3): 376–395.

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York.

Kripke, S. A. (1982). *Wittgenstein on Rules and Private Language: An elementary exposition*. Harvard University Press, Cambridge, Mass.

Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4): 343–377.

Macarthur, D. (2004). Naturalizing the Human or Humanizing Nature: Science, Nature and the Supernatural. *Erkenntnis*, 61(1): 29–51.

Macarthur, D. (2014). Subject Naturalism, Scientism and the Problem of Linguistic Meaning: Critical remarks on Price's 'naturalism without representationalism'. *Análisis: Revista de Investigación Filosófica*, 1(1): 69–85.

Macarthur, D. (2019). Liberal naturalism and the scientific image of the world. *Inquiry: An Interdisciplinary Journal of Philosophy*, 62(5): 565–585.

Maier, J. (2010). Abilities. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy* (Summer 2021 Edition), URL: https://plato.stanford.edu/archives/sum2021/entries/abilities

McHugh, C. and Way, J. (2018a). What is good reasoning? *Philosophy and Phenomenological Research*, 96: 153-174.

McHugh, C. and Way, J. (2018b). What is reasoning? *Mind*, 127(505): 167–196.

McLaughlin, P. (2000). *What Functions Explain: Functional Explanation and Self-Reproducing Systems*. Cambridge University Press.

Mercier, H. and Sperber, D. (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences*, 34(2): 57–74.

Mercier, H. and Sperber, D. (2017). *The Enigma of Reason*. Harvard University Press, Cambridge, Mass.

Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. MIT Press.

Millikan, R. G. (1989). In defense of proper functions. *Philosophy of Science*, 56: 288–302.

Millikan, R. G. (1990). Truth, Rules, Hoverflies, and the Kripke-Wittgenstein Paradox. *Philosophical Review*, 99(3): 323–53.

Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. MIT Press, Cambridge, Mass.

Millikan, R. G. (2000). *On Clear and Confused Ideas: An Essay About Substance Concepts*. Cambridge University Press, Cambridge.

Millikan, R. G. (2005). The father, the son, and the daughter: Sellars, Brandom, and Millikan. *Pragmatics and Cognition*, 13(1): 59-71.

Millikan, R. G. (2020). Neuroscience and teleosemantics. *Synthese*, doi: https://doi.org/10.1007/s11229-020-02893-9

Moosavi, P. (2019). From biological functions to natural goodness. *Philosophers' Imprint*, 19(51). URL: http://hdl.handle.net/2027/spo.3521354.0019.051

Moreno, A. and Mossio, M. (2015). *Biological Autonomy: A Philosophical and Theoretical Enquiry*. Springer.

Mossio, M., Saborido, C., and Moreno, A. (2009). An organizational account of biological functions. *British Journal for the Philosophy of Science*, 60(4): 813–841.

Neander, K. (1991). Functions as selected effects: The conceptual analyst's defense. *Philosophy of Science*, 58(2): 168–184.

Neander, K. (2017). *A Mark of the Mental: A Defence of Informational Teleosemantics*. Cambridge, Mass., MIT Press.

Neander, K. and Schulte, P. (2021). Teleological theories of mental content. In Zalta, E. N., editor, *Stanford Encyclopedia of Philosophy* (Spring 2021 Edition). URL: https://plato.stanford.edu/archives/spr2021/entries/content-teleological

Dutilh Novaes, C. (2021). *The Dialogical Roots of Deduction: Historical, Cognitive, and Philosophical Perspectives on Reasoning*. Cambridge University Press, Cambridge.

Peacocke, C. (1992). *A Study of Concepts*. MIT Press, Cambridge, Mass.

Peregrin, J. (2012). Inferentialism and the normativity of meaning. *Philosophia*, 40: 75–97.

Price, H. (2004). Naturalism without representationalism. In Caro, M. D. and Macarthur, D., editors, *Naturalism in Question*, pages 71–88. Harvard University Press.

Reinikainen, J. (2020). Meaning still not normative: On assessment and guidance. *History, Philosophy and Literary Studies*, 28(4): 510–526.

Rosen, G. (2001). Brandom on modality, normativity, and intentionality. *Philosophy and Phenomenological Research*, 63(3): 611–23.

Russell, B. (1919). *Introduction to Mathematical Philosophy*. Dover Publications.

Schlosser, G. (1998). Self-re-production and functionality. *Synthese*, 116(3): 303–354.

Schmidt, M. F., Rakoczy, H., and Tomasello, M. (2012). Young children enforce social norms selectively depending on the violator's group affiliation. *Cognition*, 124(3): 325–333.

Schmidt, M. F., Rakoczy, H., and Tomasello, M. (2019). Eighteen-month-old infants correct non-conforming actions by others. *Infancy*, 24(4): 613–635.

Sellars, W. (1953). Inference and meaning. *Mind*, 62: 313–338.

Sellars, W. (1954). Some reflections on language games. *Philosophy of Science*, 21: 204–228.

Sellars, W. (1967). Some reflections on thoughts and things. *Noûs*, 1(2): 97–121.

Sellars, W. (1969). Language as thought and as communication. *Philosophy and Phenomenological Research*, 29: 506–527.

Sellars, W. (1974). Meaning as functional classification: a perspective on the relation of syntax to semantics. *Synthese*, 27: 204–228.

Shapiro, L. (2004). Brandom on the normativity of meaning. *Philosophy and Phenomenological Research*, 68(1):141–60.

Tomasello, M. (2019). *Becoming Human*. Harvard University Press, Cambridge, Mass.

Warren, J. (2020a). Killing Kripkenstein's Monster. *Noûs*, 54(2): 257–289.

Warren, J. (2020b). *Shadows of Syntax: Revitalizing Logical and Mathematical Conventionalism*. Oxford University Press, Oxford.

Whiting, D. (2006). Between primitivism and naturalism: Brandom's theory of meaning. *Acta Analytica*, 21(3):3–22.

Whiting, D. (2007). The normativity of meaning defended. *Analysis*, 67(2): 133–140.

Whiting, D. (2009a). Is meaning fraught with ought? *Pacific Philosophical Quarterly*, 90(4): 535–555.

Whiting, D. (2009b). The normativity of meaning (steadfastly) defended: reply to Glüer and Wikforss. Unpublished manuscript.

Withrington, E. (2019). *Basic-acceptance teleosemantics*. PhD Thesis. Accessed September 9[th], 2021. URL: https://academicworks.cuny.edu/gc_etds/3080/

Wittgenstein, L. (2000). *Wittgenstein's Nachlass: The Bergen Electronic Edition 1998-2000* (cd-roms). Oxford University Press, Oxford.

Wright, L. (1973). Functions. *Philosophical Review*, 82(2): 139–168.

Yalcin, S. (2014). Semantics and metasemantics in the context of generative grammar. In Burgess, A. and Sherman, B., editors, *Metasemantics: New Essays on the Foundations of Meaning*, pages 17–54. Oxford University Press.