Distribution Agreement

In presenting this thesis or dissertation as a partial fulfillment of the requirements for an advanced degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis or dissertation in whole or in part in all forms of media, now or hereafter known, including display on the world wide web. I understand that I may select some access restrictions as part of the online submission of this thesis or dissertation. I retain all ownership rights to the copyright of the thesis or dissertation. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation.

Signature:

_____ _____

Reuben Sass Date

Practices, Types and Pragmatic Ontology

By

Reuben Sass
M.A.

Philosophy

_____

Mark Risjord
Advisor

_____

Andrew Mitchell
Committee Member

_____

Thomas Flynn
Committee Member

Accepted:

_____

Lisa A. Tedesco, Ph.D.
Dean of the James T. Laney School of Graduate Studies

_____

Date

Practices, Types and Pragmatic Ontology

By

Reuben Sass
M.A., Case Western Reserve University, 2015

Advisor: Mark Risjord, Ph.D.

An abstract of
a thesis submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Master of Arts
in Philosophy

2018

Abstract

Practices, Types, and Pragmatic Ontology

By Reuben Sass

Contemporary ontology typically involves arguments about which philosophical theories are more parsimonious or more relevant for particular philosophical questions. What is left out of such discussion is the role of social practices and norms in determining not merely our commitments to the existence particular things, but also our definition of such basic ontological concepts as types and identity criteria. This thesis focuses on the importance to ontology of social norms and practices. I draw on Brandom's inferentialism (1994), phenomenology, and Peircean pragmatism, though my arguments do not straightforwardly reproduce the arguments of these philosophical programs. The thesis likewise reflects a pragmatic reading of Quine insofar as I argue that the content of ontology should be based on whichever ontological commitments yield better results for practices of empirical research. The first chapter of the thesis argues, against Quine, that the context for determining the identity criteria for types should extend into social practices beyond the sciences, including more everyday practices. The aim of the first chapter is to articulate a criterion for determining which social practice is most relevant for defining the identity criterion for a particular type. The second chapter explores how the very concept of a type can be defined through inferences licensed by the norms of research practices, and how types can be distinguished from each other through different inferential goals. This normative conception of a type aims to rescue the very concept of types from common counterexamples and charges of triviality, while also offering a story about how particular types can be constructed. The final chapter explores how the normative conception of a type can explain progress in research practices, in addition to further discussion of the relevance of identity criteria to social practices. As a whole, the thesis aims to show that questions of social practices and norms can result in a philosophical ontology that is more informative, accounting more thoroughly for both the construction and definition of basic philosophical concepts. Such informativeness may lead to an ontology which is more relevant to a broad variety of practices of research and investigation into the world.

Practices, Types and Pragmatic Ontology

By

Reuben Sass
M.A., Case Western Reserve University, 2015

Advisor: Mark Risjord, Ph.D.

A thesis submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Master of Arts in Philosophy

2018

Acknowledgements

The author thanks Mark Risjord for extensive advice and feedback on earlier drafts, particularly in relation to the second chapter. The author likewise thanks Tom Flynn and Andrew Mitchell for providing additional advice and for agreeing to read the completed manuscript.

Table of Contents

General Introduction: Towards A Pragmatic Ontology

What is the purpose of an ontology? Why should we bother classifying the world as made up of particular *types* of things? Perhaps we should do ontology because we are interested in figuring out whether it is true that some type of thing exists. Indeed, much of analytic ontology is concerned with criteria of ontological commitment or existence conditions—what it means to say that something exists. Different criteria are used to reach different answers about which types of things exist (e.g. numbers or sets or material objects) (Thomasson 2014). Various theories of ontological commitment have come under discussion, including Quine's quantificational theory, the truthmaker theory, and most recently, the "grounding" theory (Schaffer 2008).

In recent decades, analytic metaphysicians have by and large treated ontology as a self-enclosed enterprise (Thomasson 2014), an enterprise whose success is not to be assessed through pragmatic criteria such as usefulness for empirical research. Rather, philosophers tend to answer ontological questions, and construct ontological concepts, independently of whether some answers and concepts rather than others might help researchers construct better ontologies of their own disciplines. Arguments in ontology might appeal to values such as parsimony and relevance to other philosophical disputes, rather than the advantages for explanation, description or prediction of phenomena being studied empirically.[1]

[1] That is not to say that theories of existence are not interesting in themselves, nor even that such theories may not be relevant in some cases to empirical research, or at least to the construction of classification systems or ontologies of such research. It may even be that a good case can be made for "easy ontology," or for allowing multiple, non-mutually exclusive conditions for answering existence questions (as in Thomasson 2014).

What much of analytic ontology ignores is thus the pragmatic significance of ontology—can ontology be useful for empirical research? Which ways of formulating an ontology would be most useful or informative for empirical research? What relevance does ontology have to the practices (both everyday and scientific) through which we look at or investigate the world? In one sense, the answer to the first question is trivial. If ontology is only a way of dividing up the world into types or kinds, then of course ontology in some form is useful to empirical research. Biology in particular has relied on some form of classification since Aristotle. And the demand for classification systems has only increased with the advent of genetics and the resulting research into vast numbers of genes and protein sequences (Smith and Munn 2008). Questions of ontology have also arisen in anthropology and the social sciences (Searle 2006).

What is left unanswered, and indeed largely unasked, is the question of whether *philosophy* could add much to the ontologies of particular research practices, if philosophers were inclined to take a more pragmatic view. This question does not have such an obvious answer. The ontologies of empirical research practices can be assessed according to the quality of research produced, relative to both general and discipline-specific standards. But a philosophical ontology would seek to develop better conceptions of concepts not unique to any particular research practice—the very concept of a type, for example, rather than the concept of a particular type, such as an animal type or a type of societal institution. Can philosophical approaches to such abstract questions as the meaning of a type be done in such a way as to help particular research practices better categorize, understand, or use their own discipline-specific ontologies? Can philosophy perhaps clarify the goals of particular research practices and the question of which concepts correspond to which research practices?

One approach to the pragmatic role of philosophy is suggested by an interpretation of Quine. Quine has been held responsible for inspiring the notion of ontology as its own "science," distinct from other research practices (Thomasson 2014). But it does not seem very Quinean to say that ontology should be unaccountable for, and unassessed by, its impact on the empirical disciplines. Indeed, if ontology is a science, then, given Quine's (admittedly qualified) notion of scientific holism (Quine 1975), ontology should have relevance for other branches of science, and other branches of science should have relevance for ontology. It would not be enough merely to take certain values like parsimony and explanatory adequacy and test ontological theories according to these values, independently of relevance to the empirical sciences.

Suppose that a philosophical ontology makes some very abstract conceptual commitments—e.g. to a certain definition of the concept of a type, or a certain existence condition (or meaning of the existential quantifier). Such philosophical commitments, on the pragmatic view, would not be tested for parsimony and explanatory value *in isolation* from the impact of such commitments on the theories or ontologies of more empirical disciplines. That is, the advantage of a philosophical ontology having some value—being simple or informative, for example—would not lie in the worth of that value for philosophy alone. The advantage would instead lie in the effect that philosophical commitments adhering to that value would have in *improving* empirical ontologies (e.g. of molecular biology or anthropology) constructed with such philosophical commitments in mind. Such improvement could be accounted for through a variety of basic norms, such as accuracy, predictive power, social utility, explanatory adequacy, etc. (see Kuhn 1977).

In Quinean terms, the task of philosophical ontology would be to provide the basic conceptual and logical apparatus for the ontology of any particular research practice or theory. The philosopher could specify, for example, the logical form of an existence condition without deciding what types of things do and do not exist (Quine 1948). Philosophy could explain (Quine 1975) the logical relations between the basic concepts (e.g. theoretical terms, theory formulations and observation sentences) used to construct a theory, without deciding too much about the content of a theory. The detailed content of a theory or an ontology—the decision about which particular entities exist and what their relations are—would be left to the empirical sciences. But philosophy would have a role in providing the most fundamental standards for making such decisions. The idea is that a good apparatus of philosophical ontology would stack the deck in favor of *better* empirical research, by offering insight into the structure of the conceptual and logical space within which empirical ontologies would be constructed. Such insight might allow for better ontologies by providing fundamental standards for more rigorously interpreting and codifying the information gleaned from empirical research into an ontology.

Indeed, the quantificational theory of existence advocated by Quine could be read as leaving it up to the sciences to decide what does and doesn't exist. That is to say, for something to exist is for it be the value of a variable, such that that value "has to be reckoned among the entities over which our variables range in order for our affirmations to be true" (Quine 1961, qtd. in Wetzel 2009, 25). More particularly, the class of affirmations in question are held to be the "statements affirmed in a theory" (Quine 1961, qtd. in Wetzel 2009, 25). (The values of variables can be expressed algebraically or as pronouns). The theories in question

are, for Quine, scientific theories. So existence questions can be determined by which entities, as values of variables, must be posited by the sciences themselves in order for the theoretical statements of the sciences to be true. Indeed, the standards for existence would be determined from within the scientific theory itself.[2]

What doesn't exist comes down to what the theory doesn't need to posit for the theory's statements to be true. Such unneeded (and thereby non-existent) entities include entities which could be eliminated through references to a smaller and simpler set of entities. E.g. if one could find a way to preserve the truth of all theoretical statements referring to species by just referring only to individuals, then one could eliminate species as an existent entity (Wetzel 2009, 28). Two pragmatist theses seem to follow at least from the spirit, if not the letter, of such Quinean views: that the terminology in which existence claims are made should be that of the best current scientific theory (perhaps regimented in predicate logic), not a special "philosophical" language, and that the truth of particular existence claims should be determined by the content of scientific theory, not by philosophy.

To a certain extent this pragmatic view of Quineanism is what I adopt, although I take a much broader view than Quine did of the appropriate contexts for ontology. Indeed, in the first chapter, I argue that the relevant contexts for ontology should expand beyond scientific theory

[2] Note that the quantificational theory sets an epistemic constraint on ontology (Wetzel 2009; Schaffer 2008). We are committed to whichever entities we need to *refer* to in our theoretical statements, such that those statements turn out to be true. If we cannot find a way of preserving the truth of our theoretical statements while substituting in references to a mereologically simpler class of entities (e.g. atoms or molecules *instead* of organisms), then we are committed to the existence of the more mereologically complex entities. So, on this view, our ontology is limited by our epistemic capacities, as evidenced in our theories: we must commit ourselves to whatever we don't know enough to "paraphrase away."

to incorporate a variety of research practices and other social practices. I argue that in some cases the most informative identity criteria for general terms (e.g. types) are not couched in scientific categories. That is, less scientific identity criteria sometimes pick out those sets of particulars which best serve particular practices' goals for inference or information-gathering. And an ontology of practices should determine how good an identity criterion is relative to some particular practice, based on how well that criterion fulfills that practice's goals for inference.

Moreover, I follow another tradition than Quine in emphasizing the importance of norms in delineating the basic concepts of ontology. The specifics of the framework which I present in chapter 2 owe more to Brandom (1994) and phenomenology (particularly Merleau-Ponty and the early Heidegger) than to Quine. That is, norms, contexts of use (as practices), and conditions for inference play a central role in the framework. For Quine, instead, there is not such a focus on articulating norms. The Quinean position may be that norms, standards and rules are best left to the sciences. In later work, (1975) Quine does briefly allude to the importance of different values or standards for good translation. But Quine does not offer a systematic account of the role of normative standards in constructing a better ontology.

I do agree that the specific content of norms (e.g. norms of inferring dispositions from properties) is best left to empirical research rather than philosophy. Nonetheless, philosophy can provide a general roadmap for how such norms are constructed and how such norms determine the content of types. The roadmap is like a story of the process of how and why we give content to certain types, and more particularly how such content-ascription should be done relative to our inferential goals. I take up this task in chapter 2.

So which ontologies, or ways of dividing up the world, are most useful or relevant or otherwise important for the practice of empirical research? Which ontologies help us better understand, use, or categorize such research?

I do not believe that a philosophical framework has much to say about the content of particular types of things relevant to particular research practices. But that does not mean that such a framework has no role to play. Indeed, the literature in philosophy and the empirical sciences often makes use of the distinction between a "type," an abstract object with no spatio-temporal location, which somehow refers to or models (Wetzel 2009) its particulars, and that of a token, or a spatio-temporal thing. (Chapter 2 is devoted to offering a better-articulated conception of a type).

Suppose one were to take pragmatic significance to be ontology's primary goal, contrary to the instincts of most contemporary metaphysicians. My claim is that a philosophical approach to ontology can address topics of pragmatic significance too general to be covered by any single research practice, including: an informative conception of what a type is, criteria for determining which practices should play a role in defining which types, and a schema for how such types are constructed through the norms of research practices. Through such a conception of types, one can even analyze what it means (or ought to mean) to say that research practices have made progress in finding useful information, while offering a general standard for assessing such progress.

In sum, a pragmatic ontology offers the prospect of better conceptual tools for better categorizing, understanding and using the information discovered through various research

practices. I cannot fully deliver on that prospect, but I can try to show its plausibility. Indeed, I do not claim to have constructed a pragmatic ontology which provides everything I claim such an ontology could provide. But I do want to show that a philosophical framework can, through a variety of conceptual tools, have something non-trivial to say about all the topics given above. Such conceptual tools are drawn primarily from contemporary notions of properties and dispositions, from Brandom's conception of inferences and norms, and from the emphasis on practices and contexts of use which appears in some phenomenology (particularly the early Heidegger (1996), Sartre (1992) and Merleau-Ponty (2013)). The project is broader than much philosophy of science insofar as it does not commit itself to the norms of one particular region of science, but aims for an even more general framework that can likewise apply to other social practices. Research practices still do have pride of place in determining norms of inference (and thus making conceptual content "explicit," in Brandom's (1994) famous phrase).

Chapter 1: A Criterion for Better Identifying Object Types in an Ontology of Social Practices

1. Ontology Across Social Practices: Selecting Better Contexts

In contemporary ontology, there has been concern over whether the answer to the question of "what there is" should vary. The answer might vary with what subjects do, what information subjects have, which sets of concepts and conceptual relations subjects choose, or some other such subject or observer-dependent variables. There may be big differences

between all these types of subjective variables, or "contexts."[3] Some, notably Quine (1969) and Kuhn (1996, 1977), maintain that how the ontologist *answers the question* of "what there is" should vary with context.

Indeed, the question of context in contemporary ontology can be understood through Quine, for whom the relevant type of context was that of a theory (Quine 1969, 1975). The Quinean ontologist answers the question of "what there is" from within the context of a theory (Quine 1969, 1975; Smith 2003). Though a Quinean conception of a theory may not be limited to a certain canonical set of sentences about what there is (van Inwagen 1998), in Quine's view theories at least emerge out of certain sentences concerning relations between "theoretical terms," such as "electron" and "molecule" (Quine 1975, 318). Together with "observation sentences" such as "This object is blue at such and such a time and place," a theory implies other observation sentences about what there is.

As Smith (2003) points out, Quine does not give the ontologist a means of doing ontology outside the context of whichever theory she has chosen. Quinean ontology nonetheless presupposes that the ontologist shouldn't have too much trouble selecting a theory. On a Quinean view, the ontologist could just choose the theory that contemporary scientists have good reason to believe is the best. Indeed, for Quine, the task of ontology is to specify the entities whose existence scientific theories are committed to (Quine 1975; Smith

[3] Both behaviors and theories would be contexts in this sense. To clear up some ambiguity about the use of "context," one might distinguish "evaluative" contexts and "behavioral" contexts. The theory would be the evaluative context, and the information about the behavior to be evaluated would be the behavioral context.

2003). Admittedly a chosen theory might be refined by introducing some different theoretical terms and observation sentences.

What Quinean ontology lacks are general criteria for determining which theories are better or worse for answering a given question about what there is. Kuhn (1977), by contrast, does offer some such criteria via a list of norms which cut across theories. Kuhnean norms aim to provide the scientist with good reasons to choose one theory rather than another. Such norms include, for example, consistency, simplicity, and accuracy.

Perhaps Quinean ontology omits such criteria for choosing among theories because the Quinean would not see the need for such criteria. If contemporary scientific theory is the only relevant context for ontology, then, at least on Quine's view of scientific theory, there may not be many contexts to choose among. Indeed, Quine viewed science as "a considerably integrated system of the world even now [in 1975], though the explicit reduction of major branches to theoretical physics is incomplete" (1975, 314). As a pragmatist, Quine did not subscribe to a *completely* integrated view of the different branches of science, such as physics and biology. Yet Quine was largely holistic in thinking that there was a great deal of commonality that could be achieved in scientific theory through the use of mathematics and formalization (1975, 314).

I do not mean to dispute Quine's point about scientific holism; that is not the problem on which I focus. Rather, my argument is that it is only by narrowly limiting the content of ontology to scientific theories, chiefly contemporary theories of the natural sciences, that

Quinean ontology can avoid the problem of having to provide general criteria for choosing better contexts with which to specify "what there is." But why should it be a problem to limit the content of ontology to scientific theory?

The problem is that there are indeed good reasons to choose among many contexts *other* than scientific theories in specifying "what there is." For examples of contexts more varied than contemporary scientific theories, one can turn to social practices. Sciences are social practices, but there are many social practices which are not sciences. The theory of social practices has been much discussed in recent sociological literature (Reckwitz 2002; Shove, Pantzar, and Watson 2012: Law and Lien 2012). According to Reckwitz's (2002) theory, for example, a practice is a "routinized type of behavior which consists of several interconnected elements: forms of bodily activities, forms of mental activities, things and their use, a background knowledge in the form of understanding, know-how, [and] states of emotions" (249).

If the content of ontology were to be coextensive only with scientific theory rather than the much broader array of social practices,[4] then ontology would leave out a great deal of information about such practices. Yet many fields of research, such as anthropology and sociology, are deeply concerned with social practices and behaviors that cannot easily be reduced to scientific theory. Such research fields could, and indeed already do, benefit from

[4] The use of social practices as a context for ontology is also suggested by Haugeland's (2000) reading of Heideggerian "comportments towards entities." In Haugeland's reading, those object types which exist for a given social practice depend on the uses which such practices would have for "what there is."

looking at ways to categorize the social practices they study in terms of categories of "what there is." Consider the recent "ontological turn" in anthropology and sociology (Kohn 2015; Palecek and Risjord 2013), of which I later provide an extended example. I do admit that trying to write about ontology may have left me somewhat biased about the usefulness of ontology.

Salmon farming and whitewater rafting are examples of practices which I later use to illustrate the problem with Quine's lack of a general criterion for choosing better and worse contexts in ontology. The best categories of information for object types used in those practices are not given by scientific theory, or so I argue, based on empirical literature (Law and Lien 2012). I aim to illustrate the point both that there should be a general criterion for selecting better contexts, and that such a criterion should not presuppose that scientific theory will always be the best context for evaluating behavior.

If ontology is to focus on such a broader variety of contexts, then the need for a criterion for choosing better or worse contexts becomes all the more important, since we would have so many more contexts to choose from. Indeed, without such a *normative* criterion, the ontologist of social practices might face the kind of criticism that prompted Kuhn to propose his list of norms for theory choice. To wit, an ontology of social practices without such a normative criterion might well be criticized as "a matter of mob psychology" (Kuhn 1977). I take "mob psychology" to mean something like a lack of good reasons with which one could defend one's choice to specify "what there is" in some way rather than another.

Smith (2003) raises a problem similar to that of "mob psychology" in asking, "How, ontologically, are we to treat the commitments of astrologists, or clairvoyants, or believers in Leprechauns?" Let's refer to object types identified through properties with little basis in scientific theory as "weird things." Smith suggests there is no good reason to include such weird things in our ontology, because there is no good reason to believe that such weird things have the properties attributed to them. Smith's point is that the example of such weird things undermines the case for a certain sort of ontology—an ontology concerned with a broad variety of social practices dealing with things of little use to scientific theories.

My response to Smith is that even an ontology of social practices, as I conceive it, would be formulated from the perspective of a researcher. Accordingly, the criterion for selecting better contexts which I propose in the next section adopts the researcher's perspective. Take, for example, a voodoo doll. The voodoo doll is a weird thing. Yet the ontologist explaining or describing the relevant anthropological information about a voodoo doll would not regiment such information into properties such as "having the capacity to injure one's enemies, conditional on being stuck through with a pin." Rather, anthropologically relevant properties of the voodoo doll would regiment *cultural beliefs or attitudes* conveying information about the place of such a doll in the culture being studied. Categories of cultural beliefs and attitudes can be defended as providing the information about the doll which best fulfills the anthropologist's goal of understanding how the doll is used within a particular culture. Thus I think one could have good reason to include even weird things in ontology, with the all-important caveat that

the cultural and behavioral context should be very well-specified in terms of the anthropologically relevant categories (e.g. cultural beliefs and attitudes about the weird things).

Admittedly, the approach to ontology on which I focus may seem more bottom-up than top-down. For a top-down ontology, at least, the ontologist would hope to be able to regiment entities in terms of the most general categories possible, which would typically be those of scientific theories. In a bottom up-ontology, by contrast, it may be difficult to find general categories for all the object types which would be included as relevant to the study of social practices.

Even so, my position does not suffer by admitting that certain domains of ontology, chiefly those domains dealing with scientific theory, could still be expressed in more top-down fashion in terms of scientific generalizations. I am only saying that a bottom-up ontology of social practices can be useful to researchers, so long as such an ontology of social practices contains a criterion for selecting between better or worse contexts for categorizing information about what there is.

The perspective I take is broadly that of applied ontology, as distinct from more *a priori* approaches to ontology. Applied ontology is focused mainly on specifying the categories of information useful for fulfilling the interests and goals of researchers (Smith and Klagges 2008). Indeed, I do not deal here with issues of meta-ontology—what ontological commitment means and how such commitment ought to be expressed (Quine 1948; van Inwagen 1998).

Finally, my argument agrees with Quine on one important count. I think that ontology ought to take for its content the answers which empirical research gives us to the question of what there is. My focus on such empirical research just extends beyond the perspective of natural scientists. Such an expanded perspective, I will argue, makes it all the more important to have a criterion for choosing better contexts.

2. Selecting Better Categories for Identifying Object Types for a Given Research Question: An Extended Example

The criterion which I want to develop is aimed at a particular problem. The problem is that for the same object type—say "salmon" or "river rapid"—a given way of categorizing information about that type may be good for one context but bad for another. This problem becomes more prominent when ontology expands beyond the scope of scientific theory. If, for certain social practices, defaulting to scientific theory may not always give us the best information about object types, then it should make sense to look for a more general criterion for selecting better information about object types. The evaluative contexts on which I focus, then, are not theories, but rather categories of information about types of objects.

The goal of the criterion is to find the best categories of information for uses of objects by subjects for some behavioral context. A behavioral context would be defined by a given type of use of a given type of object by a given type of subject. Such uses of objects can be associated with a broad range of social practices, including the examples which I provide below of salmon farming and rafting. In this section I will work through such examples, so as to

explain how a criterion for selecting better categories of information would be useful to the ontologist. I then articulate the criterion, which I explore further in the following section.

I believe that Law and Lien's (2012) study of salmon farming furnishes evidence in support of my criterion. I argue that such examples illustrate why it is a problem for Quinean ontology not to have a general criterion for assessing which evaluative contexts are better than others for a given behavioral use. Because Quine lacks criteria for determining which evaluative contexts are better or worse, I suggest that the Quinean has no good answer to the particular problem of how best to categorize an object type within different behavioral contexts. In sum, the ontologist would do better ontology by having a criterion which could both select among evaluative contexts not limited to scientific theories, and which could provide good reasons to choose one evaluative context over another for a given behavioral context.

Let's take the example of the salmon, the relevant object type for Law and Lien's (2012) sociological study of salmon farming. If the ontologist were a strict Quinean, how would she go about selecting categories of information for objects of the type "salmon"? First, the ontologist might recall Quine's view that the content of ontology is coextensive with the content of science. Accordingly, the strictly Quinean ontologist might look to scientific theory about salmon. Biology offers such a theory, identifying salmon, as a sub-type of animal, in terms of categories of phylogeny and morphology and genetics.

The ontologist, however, faces a difficulty which is unresolved by resorting to scientific theory. Categories of salmon handling may be better than categories of biology for a certain behavioral context, namely that of salmon farming. But without a criterion of assessment for selecting certain categories of information as better than others for a particular behavioral context, our ontologist might not be able to articulate and defend a good reason for why such categories of salmon handling may be better than categories of biology.

Consider using the categories of biology to regiment the information about salmon needed to explain or describe the salmon farmer's processing of salmon. Would such biological categories serve as well for such explanation and description as would categories of salmon handling? Law and Lien (2012) make a convincing case that the answer is "no." My question is: how can the ontologist offer good reasons to defend the convincingness of Law and Lien's answer, such that those good reasons can be generalized to other behavioral contexts as well?

Let's say the ontologist aims to select the categories of information which are best for describing and explaining the salmon farmer's behavioral use for salmon. That behavioral use is processing salmon and making them fit for consumption. Which categories of information for the object type "salmon" would best accomplish such explanation and description?

The salmon farmer uses salmon quite differently than the biologist (Law and Lien 2012). As a salmon farmer, one's use for salmon is to make those salmon fit for eventual human consumption. As a biologist, one's use for salmon would be to investigate the anatomy or

genetics of salmon. So what are the best categories of information for the ontologist to apply to the relevant object type—the salmon—for the practice of salmon farming?

Law and Lien (2012) make both negative and positive claims here. Negatively, the taxonomic and anatomic categories of biology matter little for the practice of salmon farming, since such categories do not do a good job of describing or explaining the salmon farmers' use of salmon *qua* salmon farmers—namely, farming salmon. Positively, the categories that do regiment information which explains or describes salmon farming include such properties as slipperiness and ease of grasping the salmon.

Especially given research goals of prediction or explanation, the ontologist should take account of the dispositional nature of such properties. Slipperier salmon are more difficult to grasp. That is, the property of greater slipperiness disposes salmon towards greater difficulty of being grasped. Greater slipperiness thereby helps *explain* the difficulty of grasping certain salmon. Likewise, the spatial orientation of fish and their color affect the likelihood that the salmon will be properly processed for vaccination. The properties of being incorrectly oriented and darker in color dispose salmon towards a reduced likelihood of being properly vaccinated (Law and Lien 2012, 367). So, for the context of salmon farming, categories of information about texture, handling and vaccination status would do a good job of explaining or describing *what it is* about salmon that would dispose salmon to be useable for the specified use, i.e. to be made safe and edible.

One might also offer another example here. Whitewater rafters divide rapids into seven types based on navigability. The rafters' use for rapids involves navigating those rapids in rafts. So the information which best describes rafters' navigating of such rapids can be categorized in terms of type categories such as: degree of stillness or violence in water movements, curvature of the banks and nature of possible obstructions, etc. (Gordon 2017). Such categories include dispositions towards other properties, such as the ease of navigating the rapids and the danger posed by the rapids. Thus, for example, the categories for identifying a "Class I" rapid include great "easiness" of navigation, related to "smooth water" and "gentle curves." The categories for identifying a "Class VI" rapid, by contrast, posit "extraordinary difficulty" of navigation, related to violent waters and "constant threat of death," and capable of traversal only by "paddlers of Olympic ability" (Gordon 2017).

On the other hand, categorizing river rapids in terms of flora, fauna, biomass and sediment types would not be good for the practice of whitewater rafting. Categories of flora, fauna, biomass and sediment types are more scientific, in the sense that these categories are what scientific theories of biology or geology would pay attention to. Yet such "scientific" categories would just not do a good job of explaining what it is about river rapids that disposes river rapids to be dangerous instead of easy to navigate.

What I want to claim, then, is that the better some categories of information explain or describe *what it is* that would dispose a type of object to be useable for a specified use, the better such categories would be for *identifying* that type of object within the context of the

specified use. This claim is, roughly, my criterion of assessment. I elaborate on this criterion in the following section.

The justification for this criterion of assessment, I believe, lies in a certain conception of what the content of ontology ought to be and how ontology ought to be practiced. For the conception of ontology I outlined in the introduction, ontology is about better specifying the results of empirical research about "what there is." Empirical research about "what there is" does involve explaining and describing dispositions of objects. But the researcher has to ask: dispositions towards *what*? The sociologist studying salmon farming has little reason to care about the phylogeny or taxonomy of salmon, because properties of phylogeny and taxonomy are not what dispose salmon towards being easier or harder to grasp or see or vaccinate. Yet these latter properties of ease of handling and vaccination are what salmon farming is about, because such properties determine *what it is* that *disposes* salmon to be farmable—to be used in such a way as to be made safe and edible. Accordingly, categories explaining or describing such properties are what would better identify salmon for the purpose of describing or explaining salmon farming. Such considerations lend plausibility to my criterion of assessment, as I explore further below.

3. Further Specifying a Criterion of Assessment for Criteria of Type Identity

The criterion of assessment which I propose depends on two types of uses—inferential goals and behavioral uses. It is through such uses that I aim to get more specific about the criterion of assessment which I proposed at the end of the previous section. Inferential goals

are uses for models of concepts. Inferential goals justify, explain, describe, argue about, predict, or make some other use of a set of concepts. The particular conceptual models I am concerned with include categories of information about subject and object types. Such categories provide information about what it is that disposes objects to be useable for some specified behavioral use.

Indeed, my focus here is on those inferential goals which aim to explain or describe a given behavioral use which a certain type of subject is making of a certain type of object. A behavioral use is a use that subjects have in behaving a certain way towards object tokens of a given type. Buying, for example, is a certain type of *behavioral* use that shoppers could make of a type of object, i.e. a retail product. Stealing or breaking the products would be other such behavioral uses, ones less amenable to the consumer researcher's *inferential goal*. The consumer researcher's inferential goal would likely involve explaining or describing or predicting the behavioral use of buying products. The buying of products would likely be best described or explained through information about which properties of shoppers and products would dispose shoppers to buy more or fewer products.

If the inferential goal involves, for example, prediction or explanation, the ontologist can assess how well particular models predict or explain the way certain types of subjects or objects do in fact behave. Such assessment might in turn make use of Kuhnean norms such as consistency, accuracy, and simplicity.

What I am defending here is a criterion of assessment. The criterion of assessment assesses another sort of criterion—a criterion of *type* identity. The criteria of type identity I have in mind include categories of information about subject or object types. In particular, a criterion of type identity is some finite set of ontological categories intended to regiment information about a given subject or object type. Such ontological categories could include types of dispositional properties, as well as relations to other object types or process types. The most important categories for researchers, as I have suggested, are likely to be dispositional.

Such criteria of type identity are not primarily concerned with how to tell which objects or subjects *count* as tokens of some type. That is not to say that criteria of type identity might not be used to determine which objects count as tokens of a type, but that is not the main function of such type identity criteria. The main function of such type identity criteria is to identify an object type for a given behavioral use.

To rehash Law and Lien's (2012) example, a salmon farmer can be best described as identifying tokens of the type "salmon" in terms of information which the ontologist could model through categories of texture and ease of handling. Such categories of texture and handling better explain or describe (fulfilling some inferential goal) *what it is* about salmon that *disposes* salmon to be farmable. A biologist, on the other hand, can be best described as identifying the type "salmon" in terms of categories of taxonomy, phylogeny, and genetics. Such biological categories better explain or describe *what it is* about salmon that disposes salmon to display the properties or behavior in which the biologist is likely to interested.

So now I can specify my criterion of assessment in the following terms:

For some inferential goal I(U) of some behavioral use U which some type of subject has for some type of object O, and for some criterion of identity C of O:

(1) The categories on C should provide as much relevant information as possible concerning what it is about O that disposes O towards being useable for U. C's categories are relevant insofar as such categories improve how well I(U) is fulfilled. Accordingly, the more relevant information which C provides concerning what it is about O that disposes O towards being useable for U, the better C is for identifying O within the context of U.

and

(2) For any criteria of identity of O which fulfill I(U) equally well, that identity criterion is best which has the fewest categories. This stipulation is intended as a check against superfluous categories.

The categories selected through this criterion of assessment would then constitute a better *criterion of identity* for that type of object within the context of the specified behavioral use. Keep in mind the difference between a "criterion of identity" and a criterion of assessment. The criterion of assessment would assess how good a particular criterion of

identity is for a particular behavioral use. The criteria of identity would regiment information about a given type of object for a given use.

Some might maintain, following Heil (2002, 141), that identity criteria should merely determine what counts as some sort of thing, rather than serving as "techniques for identifying." But I do not think there is good reason to say that identity criteria must only be about determining what counts as what, while somehow controlling for or holding constant our information needs in identifying something.

As Horsten (2010) points out, the degree to which identity criteria should serve as a "way of telling" (in Anscombe's phrase) depends on the things we are identifying and our purposes in identifying them. In the context of applied ontology, it does make sense for such identity criteria to conform to the ontologist's information needs. After all, applied ontology deals with types of things which researchers can study and subjects can use.

To effectively identify object types for some use, ontologists need to categorize information about what disposes those objects to be useable for the use being researched. Accordingly, the criterion of assessment I provide does not allow the ontologist to select whatever information she wishes in modeling an object type. The criterion of assessment imposes a normative commitment for selecting better categories of information.

Chapter 2: Constructing a Normative Conception of Types

1. Towards a Normative, Inferentialist Conception of Types

The literature in philosophy and cognitive psychology is replete both with references to "types" and "tokens" and to discussions of particular types and tokens. As Wetzel (2009) has pointed out, the distinction between types and tokens trades on the distinction, respectively, between abstract objects and spatiotemporal particulars. There may not be too much trouble in grasping the concept of what a spatiotemporal particular is. But the concept of a type as an "abstract object," and how such types might relate to their particulars, has proven hard to pin down. Indeed, Wetzel (2009) has exhaustively catalogued various objections to recent theories of what a type is. Many such objections have to do with the difficulties in defining types according to kinds, species, necessary and sufficient conditions, or sets of tokens.

My aim here is not to defend any of these more recent theories of types, but rather to expand and defend an old notion of what a type is. I would suggest that in trying to get clear about what types are, we should take as our starting point the suggestion that types are norms or rules. The notion of types as something like laws, rules, or norms dates back to Peirce (1931), the originator of the type-token distinction in its modern conception. A normative conception of types can also be found, at least for word types, in Sellars (1947; qtd. in Brandom 1998). Peirce maintains that a type, or "legisign," is "a general law that is a sign. This law is usually established by men. Every conventional sign is a legisign. It is not a single object, but a

general type, agreed upon" (Peirce 2.246; Wetzel 2014). Such a conception of types as norms has received rather short shrift in recent years, likely because the nature of such norms remains elusive.

Indeed, the contemporary intuition of types as abstract objects *qua* sets or categories or concepts of things reflects a certain "bias towards the substantial"—a bias towards things or representations of things. Such a bias towards the substantial may work very well for many tokens (particularly object tokens). I do not even deny that sets or categories of concepts or concepts of things are useful for getting clear about the concept of what a type itself is. In fact I do make use of sets and categories of concepts in trying to get clear about what it means to say that a type is a norm. But what I do deny is that types *are* representations or sets or categories of particular things. In attempting to define the concept of a type, the difficulties we run into may stem in large part from an over-reliance on our bias towards the substantial.

Overcoming such a bias can allow us to make distinctions between types, words and strings of words, inferences about spatiotemporal particulars (tokens) referred to by words, and subjects' goals in making inferences about spatiotemporal particulars. Such distinctions may go some ways towards demystifying the concept of a type. Or so I argue, at any rate. Indeed, in expanding on Peirce's thesis, one can posit a link between the concept of *what* a type is, why the concept of a type should be useful, and how the concept of a type can be constructed. The normative conception of types aims to explain the link between these three basic questions about types. The explanation, I argue, lies in the sense in which types are norms, and in how types can be defined through the norms of our social practices and

academic disciplines. I do not try to cast the normative conception of types I defend as an interpretation of Peirce. But such a conception does pay attention to what Peirce calls the "conventional" nature of types, the ways that types might be "established by men" (and perhaps by some women as well!).

What I present is a positive argument for why we should construe types as norms. Namely, we should construe types as norms in order to do justice to the role of types in our inferential behavior, and in order to explain how types can acquire the content that allows for such a role in our inferential behavior. My primary focus, then, is not on negative arguments about what types are *not* (see Wetzel [2009]). Nonetheless, I do try to make clear at the outset why a positive conception of types is so important. Such a positive conception, in order to be informative, must address our inferential behavior. Alternative conceptions of types fail to offer a justifiable purpose for the concept of a type. Indeed, I argue that a conception of types based on family resemblance is trivial, whereas a conception based on necessary and sufficient conditions is wrong (see Wetzel 2009).

The prominent place assigned to the role of inferences, governed by the norms of social practices, bears some affinity to Brandom (1994) and the program of inferential semantics. Indeed, the notion of norms of inference, and the dependence of such norms on our practices, can put some flesh on the Peircean schema of types as norms. For what has been missing from the Peircean schema, and those who have conceived of types as norms (Sellars 1947), is just the means to explain the links between the definition, purpose, and construction of a type—the three basic questions that a conception of types as norms should answer.

The premise of my attempted answer is that our norms of inference, and our goals in making inferences, should extend into the domain of ontology as well as semantics, informing our classification and understanding of spatiotemporal particulars. Indeed, the role of inference in ontology has recently gained currency even apart from Brandom's inferentialism. Boyd, for example, (1999) has suggested that natural kinds may be "particularly rich sources of inductive inference…" (qtd. in Bird 2007), and Bird (2007; 2011) has taken up this suggestion.

The type-as-norm, I argue, stipulates which inferences count as fulfillment of a particular inferential goal. Such inferential goals need not show much scientific ambition. What it is one's goal to infer may be little more ambitious than the taste of a fish or the mineral composition of a rock. Indeed, the aim of the normative conception of types I advocate is largely to explain such everyday inferences. What I argue is that the inferences relevant to types are inferences from properties of particulars *to* the dispositions that *ought* to be manifested by a particular in order to fulfill a certain inferential goal about that particular.

The normativity of the normative conception of types lies in determining which type one *ought* to identify a particular as being a token of. I do not try to suggest there is some fact-of-the-matter about which type, as a norm, one did in fact follow in referring to some particular. I do maintain that types offer a means of determining what one *ought* to mean in referring to some spatiotemporal particular by a word or set of words. What one ought to mean, in referring to some spatiotemporal particular by some word or words ("salmon," "rock," "hieroglyphic"), is governed by whichever type best fulfills one's inferential goals concerning those particulars.

What is critical to my argument is that types are themselves caught up in a broader structure of norms. Types are justified and given content by norms of explanation and description—norms used by those practices engaged in empirical research on all the stuff in our world that can be referred to through general terms. My goal in section 3 will be to explain how types can be constructed from referential behaviors, given certain inferential goals and certain norms of description and explanation used by research practices.

2. The Inadequacy of Alternative Conceptions of a Type

A positive conception of types—an account of what types are—should aim at explaining *what* it is in virtue of, or by which criterion, some spatiotemporal particulars count as tokens of the same type. What I argue is that if the concept of a type is to be retained, it is not enough to argue for abstract objects. One needs to offer a positive conception of types that shows why types have some explanatory or informational value. But the two main accounts of types, based on necessary and sufficient conditions and on family resemblance, both fail to offer such a positive conception. Nonetheless, dispensing with the concept of a type does not do justice to our inferential behavior. It is the explanatory role of types in our inferential behavior that can rescue the concept of types from elimination, or so I argue in later sections.

Perhaps the most obvious approach to defining a type τ is to try to posit a set of necessary and sufficient properties which spatiotemporal particulars must fulfill in order to be tokens of τ. After all, such an approach could allow for clear criteria about what it would take for a spatiotemporal particular to qualify as a token of τ. But if the notion of necessary and

sufficient conditions fails, for some types, to serve as a criterion for being a token of a given type, then a general criterion for being a token of a type cannot rely upon that notion.

Many types, however, cannot be defined in terms of sets of necessary and sufficient properties, or even some single, non-trivial property (other than being a token of the respective type) manifested in all the tokens of that type but not in other objects. Indeed, there may be some tokens of the same type which may have in common few or even no shared properties, as Wetzel (2009) points out at length. Indeed, this is why Wetzel (2009) rejects the claim that a type definition can be analyzed relative to some set of shared properties belonging to a kind. Tokens of the same type of word, for example, may often be pronounced or spelled quite differently, and may be written in very different fonts. Even tokens of object types may have few common properties. Two different coffeemaker tokens, for example, may look and operate very differently. (Consider the difference between an electronic Keurig and a glass Chemex). If one of the coffeemakers were broken, the coffeemakers might not even share the property of being able to brew coffee (or at least good coffee!).

This inability to pin down necessary and sufficient conditions also poses a well-known problem for attempts to define species types. Even if we focus on purely scientific definitions of species, we find that there is no single, generally accepted definition of a species that encompasses all the alternate definitions (Wetzel 2009; Bird 2008, 2007; Boyd 1999). Indeed, such considerations have motivated Boyd's proposal that species types be conceived of as sources of inductive inference (Bird 2008; Boyd 1999). The problem, again, lies with the fact

that as with tokens of other types of types, there is no single property shared by all members of a given species which is not shared by members of other species (Bird 2007).

In considering such difficulties of definition, we may fall back, as does Wetzel (2009), on the notion of family resemblance. But if family resemblance just means giving up on necessary *and* sufficient conditions, then we haven't made much progress towards a positive conception of types. Indeed, one might ask the question: family resemblance between what? Between tokens, between definitions of types, or between both tokens and definitions? I should hope that between tokens of the same type there is some resemblance, familial or otherwise.

The point is that if we cannot offer a positive conception of a type, then we may have little reason to retain the concept of a type at all. Lacking such a positive conception, the type-token relationship would add little in the way of explanation or information. Indeed, there should be two basic criteria for a positive conception of a type: (1) A positive conception of a type would specify the information specific to the very concept of a type rather than to particular types, so that the concept of a type would not collapse into other concepts ("species," "word," "class") with more readily definable extensions, and (2) A positive conception of a type should explain something non-trivial about the type-token relationship, or the relation between abstract objects and particulars. But family resemblance alone cannot offer such a positive conception. Or so I argue.

If one relies on family resemblance in ontology, the most obvious argument against types is that family resemblance makes the whole concept of types unnecessary. Indeed, Wittgenstein's (1953) own discussion of family resemblance may be taken as advocating the elimination of types and universals, as Wetzel (2009, 160) notes. To say that there is a family resemblance between certain particulars may just be to say that certain relations hold between such particulars. How does the concept of a type restrict or specify these relations in a way that would provide some information which one doesn't otherwise get from putative examples of types, such as species or linguistic forms? Such examples already have well-defined extensions, regardless of whatever one takes the general concept of a type to be. E.g. to say that there is a family resemblance between members of a species may just be to say that, out of the set of members of that species, there are various subsets whose elements share many of the same properties. Which properties are shared or would need to be shared, in order for particulars to count as members of that species, would simply be a matter of linguistic and scientific practice. Types would then be superfluous because types would add no information to concepts, like species, with extensions defined independently of the concept of a type itself.

Against such attempts to eliminate types, Wetzel argues that one would need the concept of a type in order to quantify over general terms such as "species" and "word." The argument is that it is very cumbersome to reduce sentences referring to plural or general terms to sentences referring to individuals (e.g. through a system such as Goodman's calculus of individuals). According to Wetzel, such reductions may be so cumbersome that there are likely some sentences incorporating general terms which one just could not reduce to sentences

about individuals. Wetzel offers the example of a complex sentence which quantifies over species and sub-species, suggesting the impossibility of reducing such a sentence to a sentence referring only to individuals.

Perhaps it is true that certain terms cannot be reduced to references to particulars.[5] Perhaps to quantify over some object, references to which cannot be reduced to references to particulars, is to quantify over an "abstract" object. Indeed, abstract objects may serve important purposes in our language and thought, even though abstract objects are not physical particulars (Asher 2012). One could even stipulate that such abstract objects are to be referred to as "types."

But such a stipulation would be unjustifiably broad. Not all abstract objects are types. After all, there are sets and classes, which may not reduce to types, as Wetzel (2009) herself argues (based on the disanalogies between species and sets). There may also be many other abstract objects in linguistics (e.g. Asher 2012). How then would one distinguish those abstract objects which are types from those which aren't?

[5] Wetzel's argument, however, may miss the difference here between the epistemic and the ontological. It may be very difficult to reduce or paraphrase certain sentences referring to general terms to certain sentences referring to individuals. But epistemic difficulty is not ontological impossibility. Whether we have the knowledge and logical system required to perform a reduction is an epistemic question. Whether it is true that a sentence S', as a reduction of another sentence S, describes the same features of the world as S—this is an ontological question. A sufficiently powerful computer program, or even a sufficiently dedicated and well-informed person, might well be able to apply some calculus of individuals, or some other logical system, so as to reduce all references to plural or general terms in even the gnarliest of sentences.

Wetzel (2009) maintains that there is a family resemblance between the tokens of a type. Not all abstract objects relate or refer to particulars between which there is a family resemblance. Perhaps only types refer to such particulars. The problem, however, is that one would then need a criterion for defining or identifying family resemblance in order to distinguish between, or at least justify distinctions between, those abstract objects which are types and those which aren't. But such a criterion seems to be what Wittgenstein was trying to make unnecessary in proposing the concept of family resemblance (Andersen 2000). Moreover, one could instigate a regress if one's criterion, for singling out those abstract objects ("types") which posit family resemblance between particulars, relied on positing family resemblance between those very abstract objects. And yet if one did have a criterion for defining family resemblance in relation to types, one would already have taken some steps towards a positive conception of types. The need for such a conception is an important premise of my argument.

Lacking such a positive conception for types, even if one posits abstract objects like species and classes and linguistic forms, one still wouldn't have explained what information is added by the concept of a type that cannot be added by the particular abstract objects themselves. One could simply say, for example, that a species is a particular abstract object, such that the relations between the members of a species are spelled out in a particular way (according to the norms of biology and biostatistics, for example). It would be redundant and uninformative to say that a species is also a type, insofar as a type is an abstract object and types posit family resemblance between their particulars. The nature of the relations between

the members of a species would be given more informatively by the concept of a species itself. What information would then be added by talk of types in terms of family resemblance? For it is the nature of the relations between the particular members of a species that would make the concept of family resemblance intelligible, not the other way around.

So why retain the concept of a type at all? My answer is that the concept of a type can help shed light on our inferential behavior vis-à-vis the particulars of the physical word. I have already explained that in order to defend the concept of a type from being reduced away, one needs to offer a positive conception of a type that doesn't depend on family resemblance or necessary and sufficient conditions. The challenge is to tell a story about our inferential behavior in which the concept of a type is non-trivial and informative above and beyond particular examples of types. Here I offer a few examples of the role of types in inferences, before trying to say more about what types are.

Suppose that Bob is a marine biologist studying salmon. In practicing marine biology, Bob focuses on studying the morphology and the genetics of salmon. Indeed, in his day job, Bob refers to salmon (both the type and its tokens) in terms of morphology and genetics. Yet one day Bob smuggles a particular salmon home, in contravention of workplace policy. None of Bob's research at work sheds any light on the taste and texture of salmon.

Bob's only interest in the salmon at work had been to study its genetic sequence. Yet Bob is nonetheless able to infer correctly that the salmon token he takes home, in virtue of

being a (fresh) salmon, will taste good (to many people, including Bob) when grilled. And it is Bob's inferential goal to identify a fish that, when cooked, will likely taste good. The practice of cooking in which Bob engages in preparing the salmon bears little resemblance to the practice of marine biology. Yet, in identifying the token as being of the salmon type, Bob can make inferences about that token which hold good for *both* the practices of marine biology and (fish) cooking.

I do not believe that the inferential benefits of type identification are limited to types of putative "natural kinds" like salmon. (Identifying a natural kind might merely allow one to draw certain kinds of inferences [Bird 2007]). Indeed, I may hear a particularly odd pronunciation of a familiar English word, or see that word inscribed in a never-before used font randomly generated by a computer. Nonetheless, my correct identification of the type of such a word token, no matter how strangely that token is inscribed or uttered, should enable me to draw other correct inferences, e.g. about the meaning of the word being inscribed or uttered (given other spoken words and a context of utterance, or given other words on a page). Wetzel (2009, 121) does take note of such inferential relations between word tokens and types, but she does not develop such relations into a positive conception of types that expands on family resemblance.

3. Linking Inferential Goals with Norms of Inference: A First Pass at Defining Types

So now it makes sense to ask: how can the normative conception of types best explain our inferential use of the concept of a type? I suggested in the first section that types are norms which stipulate which inferences fulfill a given inferential goal. Here I offer a first attempt at unpacking such a conception of types. I am then going to explain why such an attempt is incomplete, because it fails to provide a standard for *which* inferences about *which* properties matter in distinguishing one type from another. The next section details my attempt to provide such a standard.

Now types may function something like a link in a chain of (mostly inductive) inferences. One starts, for example, with certain inferential goals—certain empirical contents (e.g. properties or qualities, etc.) that one wants to be able to infer. One might, for example, want to make inferences about the taste of a fish or the mineral composition of a rock. One's inferential goals are not themselves inferences—such goals are what one wants to make inferences about (see Brandom [1994…]). But then one has to be able to determine which properties of particulars would allow one to make the inferences which it is one's goal to make.

Suppose, for example, that I want to cook a fish with a certain taste. The taste I am after is that most typically associated with salmon. So one question would be: which physical properties warrant an assertion that such and such a fish is indeed a salmon? Given such physical properties, of a certain visual appearance and texture, I can make an inductive inference that a given spatiotemporal particular is indeed a salmon. And once I have identified the particular as a token of a salmon type, I can now make certain other inferences. Namely, I

can infer that that spatiotemporal particular, being a salmon, is likely to have the taste I am after. If I am asked, for example, why I chose to buy that particular fish, I could reply, "Because it is *a* salmon, and salmon has the taste I want."

Indeed, what I am after is a means of assessing what one *ought* to mean in referring to a particular by some word or set of words (e.g. "salmon"), and also how well one has fulfilled one's inferential goals in identifying something as a token of a certain type. Note that on the normative conception of types, a token of a type is a particular for which one can make certain inferences—namely those inferences which correspond to the particular type one has selected. The content of what one ought to mean in referring to something by a particular name is governed, I propose, by the type itself as a norm. In identifying a particular as a token of a type, how well one has selected the respective type depends on how well that type fulfills one's inferential goals for the particular one is identifying. A type fulfills one's inferential goals to the extent that the inferences corresponding to that type fulfill those goals. But what then is the nature of the "correspondence" between a type and certain inferences? What is standard for determining the content of such inferences?

The conception of the type as norm towards which I am aiming is something like a function from certain properties of particulars towards certain (typically inductive) inferences that can be drawn about particulars with those properties. Take for example, some set of objects referred to (by English speakers) as "chairs," insofar as "chair" is being used with a certain inferential goal. Those objects referred to as "chairs" may be of very different shapes

and sizes. But, at least according to a type-norm that would commonly govern the use of the word "chair," one should have reason to infer that people are likely going to try to sit on the "chairs." The obvious question, then, is: how does the type governing the use of "chair" acquire its content—what properties must some particular possess in order for one to be able to draw inferences about that particular which correspond to the chair type? (And how does one provide conditions for inference which are sufficient but not necessary?). This question gets at the issue of type definition, which I am going to try to account for (in the following sections) through research practices' norms of explanation and description as applied to empirical contents such as properties and dispositions.

Such attempts at defining types so far do not pass muster, though these attempts do lay the groundwork for the definition I aim at. The problem is that, in order to explain the "correspondence" between a type and its inferences, one needs a better understanding of what an inferential goal is, and which properties the inferences corresponding to a type are supposed to be about. In particular, one should be able to determine which properties are non-trivial and relevant to the definition of a particular type. Such properties could be either the basis for inferences or the content of inferential goals. To avoid the counterexamples I discussed in the previous sections (see Wetzel 2009), one should likewise avoid laying out necessary *and* sufficient conditions that each token of a given type must satisfy.

The goal (for the ontologist) is to figure out which properties of a type would be non-trivial, constructing the type definition from such non-trivial properties. Then, the ontologist

needs to be able to make a normative determination about which sets of properties would be sufficient to allow one to (inductively) infer that some token is of a given type. And the ontologist needs to make another normative determination about which inferences would indeed need to turn out to be true for something to be considered a token of a given type. It is this latter determination that is key to defining the very concept of a type as a norm of reference justified by our interested in fulfilling certain inferential goals.

So the first question is to sort out which properties would be non-trivial to the definition of a type. Part of the issue here comes down to the definition of "property." Two objects which would seem very different in many respects may share many trivial properties. Accordingly, in contemporary metaphysics, "real" or non-trivial properties are taken to instantiate dispositions or causal features of their respective objects (Shoemaker 1980, 1998, 2004…; Bird 2007). The dispositional account of properties screens out mere-Cambridge properties, roughly defined as properties having to do with how some thing is related to other things in time and space, independently of the things' dispositions. Shoemaker provides examples of "mere-Cambridge" properties such as "being over 20 years old…being fifty miles south of a burning barn" (1980).

I don't believe, however, that dispositional properties alone are sufficient to establish non-triviality, at least insofar as type definition is concerned. All the items in my house, however different they are, share the property of being destructible by a bomb. That is, each such item, from my books to my table cloth, may be disposed to explode upon being hit with a

kiloton bomb. Nonetheless, being destroyable by a powerful bomb is probably not relevant to an account of what type of object a book or a table cloth is.

Accordingly, I propose a criterion for selecting those dispositions which are relevant to defining types. Such a criterion would be in the spirit of Peirce's conception of types as norms "established by men" (2.2…). Namely, I suggest that inferential goals themselves that determine relevance. So now I need to specify in greater detail what the content of an inferential goal is. An inferential goal is not itself an inference, but rather determines what it is that inferences ought to be about. More particularly, the criterion for inferential goals I have in mind would select dispositions based on the role of such dispositions in explaining an object's useability for particular uses.

Such a criterion ought to select those dispositions which best explain or describe how certain spatiotemporal particulars can be useable for some use to which such particulars are put. Note here that I construe "use" very broadly, as any type of behavior, from linguistic utterance to perceptual observation to manual manipulation to scientific research, in which subjects could engage in relation to spatiotemporal particulars.

The question I am considering is how the conception of inferential goals as criteria of relevance can unpack the general conception of a type as a norm. I have said that a type is a norm governing what one ought to mean, given one's inferential goals, in referring to some particular by some word or string of words, e.g. "book" or "chair." Types can thus be conceived

of as functions from properties of a particular to inferences that one can draw about what the relevant dispositions of that particular are—i.e. which dispositions of that particular will fulfill one's inferential goals. The dispositions or properties that will fulfill an inferential goal, I maintain, are those dispositions which best explain or explain how the particular is useable for the use to which one wants to put it.

4. Developing a Criterion for Defining Types through Inferential Goals

Let's return to the example of the salmon with which I started the previous section. In shopping for a fish, or some other preparable food, one's inferential goal may be to select that fish which has the taste one would want it to have. So eating is the use to which the particular is to be put; and what explains the usefulness of a particular fish for eating would be the disposition of that fish, for example, to have a certain taste and nutritional profile. Accordingly, if it is one's inferential goal to infer whether a particular is good to eat, then one would want a means of making such an inference about particulars. A certain physical appearance and texture may be associated with a certain taste and nutritional profile. (The nature and degree of this association is a subject for the norms of research practices, which I discuss below). It is not the purpose of a type to justify this association itself. I am not concerned with justifying induction per se. Rather, my goal is to figure out a criterion for *which* types could fulfill *which* inferential goals.

When I say that the type as norm can be conceived of as something like a function, I mean that the type takes certain properties as its "input," and yields as output certain inferences about the dispositions of a spatiotemporal particular. I do not mean that the type itself provides a necessary and sufficient justification for those inferences, but merely that the type plays a certain explanatory role. That is, types can explain our inferential behavior and thereby provide good reason for what we *ought* to mean by committing ourselves to certain inferential goals. The type governing the use of "salmon," for example, need not *justify* why a certain appearance *should* be associated with a certain taste and nutritional profile. Now practices do play some role in justification, as I explain in the following section.

Rather, the positing of a certain inferential goal for "salmon" explains why, for example, we regularly do make multiple inferences about dispositions (e.g. about taste and texture and nutritional profile) from a variety of tokens with a certain (salmon-like) physical appearance. And so one can provide good reason for why a given type should be most suitable for a given set of inferential goals. That is, if one has certain inferential goals, there is a certain type that *ought to* govern one's meaning in invoking a word like "salmon"—namely the type that allows for the inferences that fulfill one's inferential goals.

Indeed, one problem is to define the scope of a type. If types indeed govern inferential meanings, one might well ask: which inferences between properties correspond to some type rather than another? The concept of an inferential goal is intended to fulfill two purposes related to this question. Firstly, an inferential goal selects for the content of a type—namely

those inferences which correspond to a given type. Secondly, an inferential goal delimits the scope of a type: even if, and perhaps especially if, we construe types as norms, types can still be individuated from each other. That is, different types correspond to different inferences in virtue of fulfilling different inferential goals.

More particularly, there are two explanatory roles at issue in the question of how to define a type. First, there is the question of how to explain which inferences ought to be grouped together into a given type. What I am arguing is that a type groups together some set of inferences which are *sufficient* to fulfill a given inferential goal or set of inferential goals. It is the inferential goal that has been selected which determines the content of a type. The content of a type includes those inferences which one can draw from certain properties of particulars (e.g. properties such as appearance) to properties or dispositions whose manifestation explains why those particulars could fulfill a given use—namely that use which is the *target* of the inferential goal, or that use about which it is one's goal to make inferences.

So the question of the scope and content of a type comes down to which dispositional properties are sufficient to fulfill that use which is the target of a given inferential goal.
We might now propose, by way of a criterion, that:

1. A type-corresponding inference I corresponds to some type τ, or is included in τ, insofar as I is an inference from some set of properties of some spatiotemporal particular x *to* some dispositions of x which would dispose x to be useable for some use U. U is the

"target" of some inferential goal G so long as the relation between U and G is described by criteria (2) and (3).

2. The inferential meaning or definition of τ is given by some set of type-corresponding inferences $\{I_1 \ldots I_n\}$ if and only if $\{I_1 \ldots I_n\}$ is *sufficient* to explain why x would be useable for U. $\{I_1 \ldots I_n\}$ fulfills G insofar as $\{I_1 \ldots I_n\}$ is *sufficient* to explain why x would be useable for U.
3. Insofar as τ is a norm or rule, τ stipulates which sets of type-corresponding inferences are sufficient to explain G. I argue below that τ would thus stipulate which set or sets of type-corresponding inferences one ought to mean, given one's commitment to G, in using some word or words to refer to x.

The selection of a particular inferential goal then explains why certain inferences should be grouped together into a given type. It is the role of the *relevant* dispositions to explain why some particulars should be useable for a use that is the target of an inferential goal. (Recall that my use of "use" encompasses any kind of behavior in which a subject might engage relative to spatiotemporal particulars, including observation, perception, and theorizing, as well as manual manipulation).

The presence of a certain disposition D may be one factor in explaining why a particular x should be useable for U; but that does not mean that D is *sufficient* to explain why x is useable for U. Consider, for example, the type of paper which would be disposed to receive an A from a competent instructor, or to which a competent instructor would be disposed to give an A. Clarity of writing may be one important factor in explaining why a paper would be disposed to

receive an A, but clarity of writing may not be sufficient to explain why a paper would receive an A. Only in combination with certain other virtues, such as coherence of argumentation or organization, might clarity of writing be sufficient to explain why a paper would receive an A. Or, to draw on my go-to example, a certain combination of fat and protein may be one factor in explaining why a certain fish should be good to eat, but having a certain nutritional profile may not be sufficient to explain why a fish should be good to eat. Other factors, such as having a tolerable taste and being fresh, might also be important in explaining why a fish would (be disposed to be) good to eat.

So on the normative conception of types, a given type explains certain patterns of inferential behavior. A given type answers the question: which properties are sufficient to *explain*, when clustered together, the fulfillment of a particular inferential goal? But a type answers that question by stipulation. It would be hopelessly circular to say that a given type explains the very norms by which that type is constructed. Without succumbing to such circularity, particular types cannot answer the questions of *why* certain properties should be sufficient to explain a particular inferential goal, and which inferences from properties to dispositions count as correct. Rather, I shall argue below that it is the norms of research practices which can answer these latter questions, thereby mapping out the process by which particular types can be constructed.

5. Inferential Goals, Practices and Research: A Schema for Constructing Types

So far I have tried to say what types are. But I have not explained how a definition of a particular type can be constructed relative to the norms of practices and research practices.[6] Indeed, I have been contending that a significant advantage of conceiving of types as norms is that such a conception is informative. One aspect of such informativeness involves, as I have argued, the ability to offer a positive definition, or at least a schema, of what types are. The other aspect of such informativeness is the ability to account for how the definition or schema for a particular type could be constructed. It is the latter aspect of informativeness on which I now want to focus. Indeed, by discussing the construction of types, I aim to show that one can further illustrate the criteria of type definition which I have offered.

So, to begin the process of type definition, the ontologist starts with a certain inferential goal or goals that are of sufficient importance to warrant type definition. (The notion of sufficient important is itself a normative determination, which I explore in the final section). The ontologist then ought to look at the inferential behavior of individuals with the inferential goal that the ontologist has selected.

[6] One potential problem is how to delimit the scope of a given social practice. If social practices are defined through a finite set of uses or behaviors, there may be a Kripkensteinian problem of how to delimit such a finite set given that a practice may exemplify infinite behaviors. I would suggest a pragmatic approach, whereby how good a definition of a practice is depends on how useful that definition is in fulfilling a certain research question. A practice may be a certain set of "similar" behaviors, with such similarity (and thus the scope of the practice) being determined based on how well the fulfillment of a certain set of behaviors predicts or explains the fulfillment of the behavior being researched. The practice, I suggest, is the "master use" whose fulfillment (e.g. farming salmon) is predicted or explained by the fulfillment of the similar instrumental uses (e.g. cleaning, vaccinating, and processing salmon). Admittedly I do not have sufficient space here to argue in detail for this conception of the scope of a practice.

Now in referring to some spatiotemporal particulars, we language users use some word or set of words. The starting question for type definition is: in using some word or set of words to refer to particulars, which inferential goal are individuals committed to when they refer to such particulars?

Now there may be no necessary and sufficient fact about mental states to indicate that individuals have a certain inferential goal. Indeed, I do not mean to challenge Kripke's (1980) famous argument against a fact-of-the-matter for following one rule rather than another. What I would rather suggest, following Brandom (1994, 236), is that having such a goal would be a matter of commitment—do individuals recognize themselves as committed to having such a goal, such that they would acknowledge that they are committed to such a goal in referring to certain particulars?

Once the inferential goal, the word or words being used relative to that goal, and the particulars being referred by to by that word have been specified, the norms of practices and research practices come into play. A practice might be more or less specialized. Insofar as one is determining which type or types the practitioners ought to be committed to, one is engaging at the very least in reflection, and quite possibly in a research practice. Now the most systematic forms of research practices would correspond to academic disciplines such as sociology, biology, physics, geology, etc. But the notion of a research practice could extend beyond such academic disciplines into any practice concerned with explaining and describing human behavior or natural phenomena according to certain norms or standards. Obviously

some research practices, often but not always those corresponding to academic disciplines, would be more consistent and accurate in their explanations, and so perhaps ought to be considered as the last word, or at least the best, when it comes to type definition.

In applying the norms of a research practice, the ontologist would have two considerations about sufficiency: which dispositions are sufficient to explain the fulfillment of the inferential goal which is selected, and which properties are sufficient to infer those dispositions? Both of these considerations, I contend, would depend on the norms of the research practices as applied to the particulars which are being referred to by the type (again, holding constant the inferential goal).

It is the norms of the research practices which would deal with two problems: determining both the correctness and relevance of inferences. Regarding the issue of relevance, as I discussed in section [3?], an indefinite number of both dispositional and mere-Cambridge properties could be defined for a given particular. It is the inferential goal which narrows down and determines the relevance of properties: for the particulars to which one is referring *while* committed to a certain inferential goal, those properties are relevant which allow one to infer (the instantiation of) those dispositions of the particulars which explain the inferential goal.

The question of sufficient conditions for inference also bears upon the issue of correctness. Indeed, it is one thing to ask which dispositions are sufficient to explain a given

inferential goal, and which property or properties are sufficient to infer certain dispositions. But it is another thing to ask what the *standards* should be for determining such sufficiency. Accordingly, the research practices themselves, in applying such standards of sufficiency to research on particulars, would thereby determine which specific dispositions and properties would achieve particular inferential goals. Those inferences are correct which would be judged, given the standards of research practices and some particulars on which to conduct research, to meet such standards.

Take, for example, the case of salmon. Which texture, color, appearance, etc. would be sufficient to infer that the particular referred to (is likely to have) certain dispositions associated with particulars referred to as "salmon"? I.e. what counts as a correct inference from certain properties of color and texture and physical appearance to certain dispositions to taste a certain way when cooked, or to manifest certain types of proteins or omega-3s? And what dispositions would in turn fulfill the goal of selecting something that can be a staple of a healthy diet?

Note that the standards for what counts as a correct inference may themselves be context-sensitive. In most jungles on earth one may be able to infer that a brightly colored frog is probably poisonous, but that need not be the case for every possible environment or planet on which frogs might exist. (See Brandom 1994). Even so, the norms of practices do and ought to take such environmental contexts into account in determining what counts as a correct inference. So I do not believe my conception of type definition would be defeated by the fact

that certain inferences might be correct in one environmental context but not in another. Type definitions could take into account such variation in environment or context, and norms of research practices might impose stricter or different requirements for inferring dispositions depending on environmental variation.

Now I am not going to delve into the relevant statistical issues, or try to offer a fine-grained analysis of how correctness might be cashed out in terms of sufficient reliability. My concern is merely to map out the schema for how particular type definitions would be constructed through practices and research practices, since it is my goal to defend the normative conception of types as being informative. In other words, I want to outline the steps of the process of type definition in order to show the plausibility and informativeness of the normative conception of types.

There are two types of sets which the process of type definition yields for a given inferential goal and a given set of particulars. There is, on the one hand, a set S_p of properties sufficient to infer the dispositions which are sufficient to explain the fulfillment of the inferential goal, and on the other the set S_d of such dispositions. A type definition (see above) is itself given by all the inferences that one can draw from the properties in S_p to those in S_d. Note that, for a given inferential goal, there may be multiple sets S_p and S_d.

Indeed, the notion of there being multiple sets of properties and dispositions sufficient to fulfill an inferential goal addresses the difficulties in defining types through necessary *and*

sufficient conditions. As discussed in the second section, for many types there may be no (dispositional) property or set of properties present in the tokens of one type but not in tokens of other types (Wetzel 2009; Bird 2007; Bird and… 2011). The notion of multiple S_p and S_d sets for a given inferential goal can show how type definitions are constructed through practices, while at the same time avoiding the claim that sufficient conditions for type definition would also have to be necessary ones.

The question for type definition, then, is how to construct property and disposition sets by applying the norms of research practices to some set of particulars, given a particular inferential goal. Consider applying the relevance criterion to the practice of office work, so as to construct a type definition for particulars referred to as "chairs" by the office workers. The relevant research practices involved in constructing such a type definition might draw on certain methods of qualitative research and surveying, plus certain principles of biomechanics. Suppose, then, the relevant use—the target of the inferential goal-- would be sitting upon certain particulars ("chairs") while doing one's work at a desk. The office workers themselves might readily admit that their use for "chairs" would be sitting. (That does not preclude other uses of things referred to as chairs). The question for type definition is to figure out which dispositions explain or describe why a particular thing is capable of being used for (prolonged) sitting by an office worker, and which properties are sufficient to infer that a thing has such dispositions.

Accordingly, using qualitative methods of observation and applying certain basic principles of physics and biomechanics, the researchers would attempt to discover what it is about the "chairs" that disposed those "chairs" to be useable for sitting at a desk. The sociologists might discover many such dispositions, cataloguable in some set: {being capable of bearing certain weights, allowing for a minimum standard of comfort, putting the user in a position of being able to reach the desk…}. Relevant properties for inferring such dispositions might include: {presenting a flat surface, presenting a surface of above a certain length and width, resting upon legs, resting upon a pedestal…}. Each particular chair might not manifest all the elements in either the full disposition set or the full property set. Rather, the researcher would aim to determine which subsets of properties and dispositions, taken from the full sets of properties and dispositions relevant to the inferential goal, would be sufficient to infer that a given particular would in fact be useable for sitting.

One further reason why it is informative to conceive of types as norms is that such a conception can illustrate the relation between our linguistic practices and research practices, allowing one to map responsible referential uses of language in relation to our inferential uses. That is, one might be held responsible for using the *wrong* word to refer to particulars, if one is committed to a certain inferential goal but one uses a word whose association with that goal is not licensed by the norms of linguistic practices.

In speaking of an "association" between an inferential goal and a word, I am referring to the question of which particular word it is acceptable to use in referring to some

spatiotemporal particular x, given a *commitment* to a particular inferential goal about x. The correctness of such goal-word associations would be assessed differently than the correctness of the inferences that contribute to type definition. The correctness of a goal-word association would be determined by the norms of *linguistic* practices, or the norms governing how speakers use a language and engage in speech acts within particular environments (see Brandom 1994). The correctness of inferences from properties to dispositions, however, would be assessed not in relation to the norms of linguistic practices in general, but rather in relation to the norms of *research* practices, including academic disciplines such as biology, physics, sociology, etc. Or so I have argued in section VI.

Consider, for example, if an English speaker were, in conversation with other English speakers, to use the word "book" to refer to particulars about which the speaker is committed to the goal of inferring dispositions to taste like fish when cooked. Quite apart from any wrong inferences the speaker might make about the particular "book", the speaker would be guilty of misusing the word "book" to refer to particulars. Such guilt would be assessed as a failure to adhere to the norms of English speakers' linguistic practices, which determine the inferential goals which could correctly be associated with particulars referred to by "book." If, however, the speaker then made many incorrect inferences about the nature of the "book," her errors would not be errors of linguistic practice alone, but rather errors of a more ontological sort, concerning the nature of the physical particulars which linguistic practices license one to refer to as "books."

6. Objections to Conceiving of Types as Norms

I have so far been trying to present a positive argument in favor of types as norms. But there are certain important objections, which arise both from my argument and from the literature (e.g. Wetzel 2009). If my conception of types is plausible, that conception should hold up to objections such as the ones I discuss below.

In response to my account of type definition, one might object that many of us seem to use the concept of a type without being familiar with the norms of research practices. Most of us are not researchers or scientists. Yet most of us use type terminology ("salmon, "book," "sentence"), and most of us do make the sorts of inferences which constitute what I have called a "type definition." I have, however, been arguing that how a type acquires its content—the set of inferences constituting the definition of a type—depends on the norms of research practices. How, then, can types, as concepts so seemingly ubiquitous in our linguistic practices, be defined through the norms of research practices, when so few of us practice research? Moreover, how could one be held responsible for misusing a type, whether by making an incorrect inference about a particular or using a word that, according to the norms of linguistic practices, is not associated with the inferences one wants to make about that particular?

First, I want to point out that research practices need not be entirely co-extensive with academic disciplines. Research practices, as I use the term, are practices that seek to explain or describe other practices. Academic practices may often accomplish such explanation or description better than non-academic practices, but this need not always be the case. Nor need

one actually be a researcher in order to be familiar with the findings of research from the social and natural sciences. Many students, after all, are made familiar with scientific norms, and the norms of research practices, without being expected to do such research.

More particularly to my argument, I want to re-iterate that the conception of a type I have been defending serves an explanatory role. That is, a type is a way of explaining which inferences ought to be grouped together—namely those inferences which fulfill a common inferential goal. It is this group of inferences about particulars that then becomes, through norms of linguistic practices, associated with the use of a word or set of words.

The point is that one need not understand all the complex norms of practices in order to be responsible for making certain inferences correctly, or for using words correctly to refer to particulars about which one can make such inferences. One need not, in other words, be familiar with all the inferences about particulars that could be grouped into a type in order to make some of those inferences correctly. Nor need one be familiar with the full scientific or sociological justifications for various inferences about particulars in order to correctly make those inferences and to be held responsible for making those inferences correctly. I need not be familiar with the genetics or morphology of salmon in order to make inferences about particulars which have certain properties of texture and coloration and taste.

Now I also want to address Wetzel's argument against the conception of types as norms. Indeed, Wetzel (2009, 2014) rejects Peirce's conception of types as norms because of

Peirce's requirement that every type have tokens, and she rejects the broader notion of types as norms or habits because of the problem of undiscovered types. I would first point out that even if Peirce is wrong about whether there are types without tokens, I don't see why there being types without tokens should pose a more general problem for the conception of types as norms. So I focus instead on whether undiscovered types pose a problem for the normative conception of types. Undiscovered types have instantiated tokens, but humans would not be aware that these tokens are of a new type.

Why might the possibility of undiscovered types be a problem for the conception of types as norms? Well, if types are indeed norms or rules or even habits (…), then let us suppose that such norms or rules have to be articulated and used by people. Let's also suppose that there are types of unexpressed or uninscribed sentences and presumably types of undiscovered species as well (Wetzel 2009). Such undiscovered types would either be instantiated in tokens or could be so instantiated, and yet we wouldn't (by definition) be aware of such instantiation. But types are norms, and such norms would have to be defined or articulated by people. Types would need to be articulated or defined in order to be instantiated (at least, I suppose, on a constructivist view of norms). Undiscovered types have not been articulated or defined—tokens of such types would have to be discovered as being of a new type in order for such a type to be articulated and defined. And so undiscovered types would remain uninstantiated if types are indeed norms. But undiscovered types, such as undiscovered species, are already instantiated in tokens. Thus the normative conception of types would seem to fall prey to a *reductio*.

Even apart from the possibility of realism about norms (which I don't defend here), however, the attempt at such a *reductio* misses the notion of normative structure. Individual types may not be defined if the relevant tokens—about which one could make the inferences definitive of the type—have not been discovered. Yet the same normative structure used to define discovered types would also be used to define undiscovered or uninstantiated types. Such a normative structure can be found in the norms of research practices governing how types would be and should be defined. Such norms of practices would (at least on a Brandom-style conception) determine what counts as a correct inference from properties to dispositions, and from dispositions to fulfillment of uses. Research practices (e.g. chemistry and biology and linguistics) also have conventions of nomenclature for both instantiated and uninstantiated types (e.g. the IUPAC system of naming chemicals, or biological taxonomy, or the whole apparatus of linguistic terminology and grammar for classifying sentence and word types…).

It is not, then, as if the normative conception of types would require new types to suddenly be instantiated *ex nihilo* by the discovery of new tokens, e.g. of a new species. Undiscovered types would, to use another Brandom-style term, be implicit in the norms of research practices. Indeed, it is misleading to speak of the normative conception of types as requiring types to be *instantiated* through new norms or rules. On the normative conception, a new type would merely be a new application of the same standards of type definition. The act of defining a new type, as a norm specifying correct inferences, does not instantiate that type in the sense of bringing that type into existence.

Rather, my claim is that on the normative conception of types, the existence of a type is not dependent on whether we have bothered to define that type, but rather on whether such a type *could* be defined given the normative structure of research practices as applied to both discovered and undiscovered, but discoverable, empirical findings. Defining a new type is more akin to a decision about whether and how to apply the same old (but revisable) research standards to new empirical findings. Defining a new type merely makes explicit the extension of the same normative structure to new empirical findings, or, given a new inferential goal, to previous empirical findings as well. It is in this sense that one can refer to undiscovered types as already implicit in the norms of practices.

If, for example, a new species or chemical were discovered, the way that species would be defined, and even the name of that species or chemical, would be given by the norms of such research practices. Such higher-order norms of type definition and nomenclature would be consistent across those types defined by the same research practices or set of practices. Indeed, the normative structure of practices may determine whether a new type even should be defined, given the discovery or instantiation of new tokens.

Another major objection lies in the view that, at first blush, the normative conception of types seems to be committed against essentialism. After all, I have argued that types are constructed through the norms of social practices. I do not say that types are graven in stone. But if certain clusters of certain properties and dispositions are indeed essential—being the

case in every possible world—then skepticism about the inferential value of types might be more firmly laid to rest. Indeed, skepticism about induction might lead to skepticism about the whole notion of types as norms. The inductive skeptic might claim that the inferential value of types presumes regular associations between properties and dispositions, but that reliance on the continuance of such regularity cannot be justified. Yet, in spite of putting all my emphasis on the inferential value of types, I seem to have deprived myself of a resource—namely essentialism about dispositional properties—that might serve to defend the inferential value of types against the inductive skeptic.

But I am not arguing that particulars cannot have essences. I am not denying--though I am not arguing--that at least some of the particulars referred to as "water" are indeed H_2O in this world and every possible world, and that H_2O has the same dispositional properties in every possible world. My argument does not entail otherwise. The normative notion of types only holds that how we group together inferences about property-disposition clusters depends on what it is our goal to infer. Admittedly the norms of research practices play a role in determining the correctness of inferences from properties to dispositions. The determination of such correctness, however, may simply be a case of figuring out how to best model whichever property-disposition clusters exist in reality. Essentialism about property-disposition clusters might indeed lend greater support to my arguments about being able to infer dispositions from certain properties. Nonetheless, the truth of such essentialism is still not critical to my argument. What is critical to my argument is that practices and research practices can define

standards for correct inferences from properties to dispositions, and that such inferences are stipulated by types for a given inferential goal. [7]

Finally, one might argue that my emphasis on inferential goals would place few limits on the proliferation of types, leading to a hopelessly bloated ontology. After all, can't people have any number of bizarre and useless inferential goals? Wouldn't such bizarre and useless goals, on my argument, then lead to bizarre and useless types? At the very least, wouldn't relativity to inferential goals lead to many trivial types used to refer to the same particular?

Norms of practices, however, can also govern the social utility and importance of particular inferential goals. That is, not every inferential goal would be sufficiently important to warrant the construction of a type corresponding to that goal. The question of which inferential goals are of sufficient importance would be determined through the norms of research practices concerned with explaining the uses that would be the targets of the inferential goals. (Recall that inferential goals are concerned with inferring dispositions so as to explain the fulfillment of particular uses). I do not think such norms of sufficient importance can be determined *in concreto* without detailed understanding of the relevant practices. Schematically, norms of sufficient importance ought to take into account linguistic, social, and scientific usefulness—the importance of particular types and inferential goals for understanding

[7] I suppose I am committed to rejecting the application of rigid designators to general terms. That is, I am not claiming what might be entailed in extending rigid designators to general terms: that a general term somehow refers, at any given time, to one and only one set of particulars in each possible world. But even Kripke's sympathizers (Schwartz 2002) have expressed skepticism about the extension of rigid designators to general terms, particularly given the lack of a well-defined semantics for extending rigid designators to general terms.

our practices of science and social science and language use in general. Only inferential goals that pass some threshold of importance for understanding practices ought to be included in an ontology.

Also bear in mind that the norms of research practices would govern the correctness of inferences from properties to dispositions. Thus inferential goals which cannot be fulfilled, at least according to available theory and evidence, would not be permitted in type definition. That is, if it cannot be established that there are dispositions which would fulfill a particular use selected as the target of an inferential goal, then one couldn't support the positing of a type that would correspond to such an inferential goal. The inferential goal of determining what would dispose pigs to fly around in the air on their own power would, for example, be unfulfillable, and so should not matter for type definition.

Chapter 3: The Concept of Progress in Research and the Relevance of Types to Practices

1. Introduction: The Effectiveness of Types

One can build on the normative conception to explain what's at stake conceptually in the very notion of progress in research practices. It would seem quite trivial to say that the goal of research practices is to expand and improve society's collective information or knowledge base. It might even be pretty close to an analytic truth to say so. But even if the basic goal of research is so obvious, that doesn't mean it's so obvious what we mean—or, more importantly,

what we ought to mean—by "progress" in the information-gathering work of research practices. So what should we mean, in the most general sense, by such progress? Generality here indicates a focus not on how progress is made for specific practices, or even for scientific theory, but rather how progress in society's information base is made *through* any social practice committed to research. I am not asking here about the societal conditions conducive to progress in research, though such conditions are of great pragmatic importance (Kuhn 1962, 1970). I am rather looking for a general means of assessing what constitutes progress in the information base. I argue that such progress can be assessed through a relation between the effectiveness and relevance of the types constructed by research practices.

Given the variety of goals and even of norms across research practices, can one speak meaningfully of a single standard, however abstract, for assessing progress across those practices? After Kuhn, in particular, such a question may assume greater importance philosophically. For Kuhn (1977, 1970) did not rule out the notion of progress, but rather argued that progress must be relative, or assessable only in relation to particular values or norms which research practices commit themselves to. The framework I defend is intended to argue that such a standard for progress is indeed both possible and non-trivial, given the notion of fulfilling inferential goals. That is, such a standard can apply across all research practices, and can measure progress, while still allowing for a notion of progress which is determined relative to inferential goals and norms for assessing the fulfillment of such goals.

But how does the normative conception of types, or the connection between types and inferential goals, help explain such progress? There are two elements to such progress: the

effectiveness and relevance of inferences. My first goal is to explain effectiveness, and then I take up relevance, and the connection between relevance and effectiveness.

So far I have defended the normative conception of types based on largely philosophical considerations. A pragmatist, however, might want to ask whether the normative conception of types might be of some use to the research practices themselves. I would suggest that the normative conception can indeed be of some use, insofar as such a conception would allow for getting clearer about both the goal of *better* type definition and the means by which such a goal could be achieved. [8]

Yet while I do not want to posit any types not constructed through norms, I do want to argue for the importance and intelligibility of the notion of an "ideal" type. Ideal types may only be ideal in virtue of having for their content inferences which do the best possible job of meeting certain basic norms, e.g. of explanatory adequacy or predictive power or accuracy or reliability. These basic norms may in turn be "constructed," and may vary both in content and importance with progress in research practices themselves. Yet these basic norms, notably discussed by Kuhn (1977), are often widely shared between research practices, and such norms are likely more stable and enduring than the more detailed methodological norms of research practices taken individually.

[8] The question of *better* type definition may also have some philosophical relevance, insofar as this question helps get clearer about the relation of my conception of types to common notions of realism and anti-realism. The normative conception of types may seem anti-realist insofar as that conception posits that types are constructed rather than found fully-formed in the wild.

What I want to suggest is that progress in research practices can be conceptualized through approximation to ideal types, which would do the best possible job of fulfilling those basic norms widely shared among research practices. Below I unpack the notion of ideal types in terms of the normative conception of types I have defended. Before diving into the weeds, I should explain that my approach shares in common with (scientific) realism (see e.g. Smith 2003) the notion of trying to approximate towards some set of ideal types which are taken as fundamental targets of research. In realism such types are both truly predicable of things in the external world and important for explaining the world. Ideal types are targets insofar as such types are what the researcher aims to approximate towards. But in contrast with stronger variants of such realism (Smith 2003; Bird 2007), ideal types on my conception need not presuppose that ideal types best capture features of a single world characterized by unchanging laws of nature and necessary property-disposition clusters. Ideal types need only be stable enough that they set a goalpost which is not always changing and does not change without good reason. The stability of the goalpost—as something towards which research can progressively approximate—would be due to the stability (e.g. quantitative accuracy or predictive power) of those basic norms which the ideal type would do the best possible job of fulfilling, as well as the stability of practices' inferential goals, as I discuss below. Such an approach does not rule out shifts in what would count as an ideal type for a given inferential goal (e.g. as the result of something like Kuhnean paradigm shifts in scientific theory). But such an approach should at least ensure that what counts as an ideal type for a research practice would not shift too easily or capriciously.

To unpack what it means to make progress towards an ideal type, consider first the question of what's at stake in *better* type definition. The basic idea of what makes type definition better is that certain inferences about dispositions can do a better or worse job of fulfilling an inferential goal corresponding to a type. Recall that the normative conception of types offers a way of constructing types by applying the norms of research practices to real-world behaviors. Such norms determine the inferences that one is allowed to make about the dispositions of whichever particulars one is trying to classify as tokens of a given type. Then the selection of an inferential goal determines which inferences about dispositions correspond to the type constructed through such an inferential goal. Indeed, types are defined through inferential goals, as I make clear in my criterion of type definition in section 4 of the type definition chapter.[9]

There are, for example, basic norms of accuracy and predictive power which research practices would aim to fulfill in licensing inferences. But the question is which inferences about dispositions better fulfill the inferential goal? The answer to this question determines the content of the type corresponding to the inferential goal, or the inferences which "correspond" to a given type.

[9] Some confusion may be caused by my use of the term "sufficient" in my discussion of type definition in chapter 2. I have said that types draw on those properties which are sufficient to infer those dispositions which fulfill the target use. But if such properties are sufficient for type definition, why should one bother with "ideal types"? As I use the term, however, "sufficient" is not synonymous with "best possible" or "ideal." Sufficiency would be cashed out pragmatically, as what is good enough to predict or explain the relevant dispositions for a certain stage of research or a particular research project. At least in my use, then, sufficiency does not mean that one has identified those properties which *best* fulfill the inferential goal which defines the research practice itself (e.g. of molecular biology or consumer research)

Recall here that uses are what it is one's goal to make inferences about. Uses are thereby what I call the "targets" of inferential goals (in the previous chapter). Even something so basic as a perceptual use—like having the goal of inferring what something's color pattern would look like in sunlight—counts as a "use," according to my use of "use."

So the inferential goal lies in determining which properties dispose towards fulfillment of the targeted use. But in terms of determining which types are better, there is the question of which types have inferential content that allows for more accurate prediction (or better explanation of, or some other inferential goal, etc.) of *which properties* would dispose towards the fulfillment of the targeted use. So certain types are better than others relative to an inferential goal. If there is such a thing as an ideal type, or a set of equally good ideal types, then such types would be those which would best fulfill a given inferential goal.

In reality, the research practices are often concerned with approximating towards such ideal types for the inferential goal with which they are concerned. That is, if it is one's inferential goal to predict which properties would dispose towards the fulfillment of some use *U*, then the goal of practices is often not only to figure out which properties are better and worse for *U*-ing, but often to make ongoing progress towards identifying properties which are better and better for U-ing. This progress is then a goal of approximating towards an ideal type—identifying which properties allow one to infer that particulars with such properties will better fulfill *U*.

Approximation towards an ideal type thereby involves the construction of inferences which do a better and better job of fulfilling the inferential goal towards which that type

corresponds. (Note that inferential goals may themselves be more or less socially useful or scientifically important, but that is a separate question from how well a given inferential goal is fulfilled). In the case of type definition, one may ask the question of the extent to which certain properties dispose towards fulfillment of a given use. It is the question of the extent of fulfillment that can be unpacked and assessed through basic norms of accuracy, predictive power, etc. To what extent do certain properties of some particulars $\{X_{1...}X_n\}$ predict, or explain, or otherwise account for the dispositions of $\{X_{1...}X_n\}$ to fulfill the targeted use?[10] How reliable or adequate or reproducible are such predictions and explanations? One might even consider unpacking this question of extent in terms of varieties of alethic or epistemic modality, if one were to construe prediction and explanation, for example, as different epistemic modes. (Indeed, the "inferential goals" to which I refer in the first section and again below may just be subset of epistemic modes). I do not follow up on this notion of epistemic modality here, but epistemic modality might at least serve as a clarifying device for unpacking phrases like "the *extent to which X is disposed* towards fulfilling U."

2. The Distinction Between the Relevance and the Effectiveness of Different Types

In the prior section I focused on defining what it means to more effectively fulfill an inferential goal. But in defining particular types through inferential goals, there is an important distinction between the *relevance* of particular (type-defining) inferences from

properties to dispositions and the *effectiveness* of sets of such inferences (i.e. the *content* of types) in fulfilling inferential goals. That is, we should distinguish between *whether* an inference fulfills some inferential goal and *the extent* to which an inference fulfills an inferential goal. By itself, this distinction may well be trivial. But I am proposing a pragmatic link between the relevance and the effectiveness of type-defining inferences. Such a pragmatic link may be realizable through a certain schema or procedure. This link is pragmatic insofar as it bears upon consequences for action, in this case the progress of research practices. That is, a procedure or schema for identifying more relevant inferences may in turn promote the progress of research practices in their pursuit of their inferential goals.

The distinction between the relevance and effectiveness of inferences is important because both terms are related to each other in a (pragmatically) important way, but the nature of this relation cannot become clear until one makes the distinction. Greater relevance of inferences may not always entail greater effectiveness, though there may often be an association between relevance and effectiveness. But the pragmatic claim is that a *schema* for better determining relevant inferences will also help identify more effective *sets* of such inferences (type-content), and more effective sets of inferences will in turn allow the identification of additional relevant inferences, which would in turn allow for the construction of more effective sets of such inferences—and so on, perhaps *ad infinitum*.

This reciprocally and progressively promotive relation between relevance and effectiveness is what I refer to as "virtuous circularity." Better determination of relevance "promotes" improvements in effectiveness, and vice-versa. If the thesis is justified, then a basic framework for identifying relevance should be important.

The goal is for the generality of the schema to hold across a variety of research practices, particularly in the social and biological research. To borrow Brandom's famous terms, the schema aims to make explicit what is already implicit in the progress of such practices. By making it explicit, however, one may gain an additional resource for modeling how practices do achieve progress, and for explaining what's at stake conceptually in the very notion of progress in research practices. That is, a schema for which inferences are relevant to which inferential goals may help one construct and identify more effective inferences. And the identification of more effective inferences may in turn improve the relevance of one's inferences.

In the broadest terms, what I am trying to do in mapping out the process for defining better and better criteria of identity is to defend the pragmatic significance of the normative conception of types. Such a conception aims to suggest the broadest conceptual outlines of how progress is made in research practices. The goal is to keep working towards a better and better fulfillment of an inferential goal through identifying more relevant and more effective types—so as to better approximate the ideal type for each inferential goal.

What I refer to as a model of a practice includes those properties of subjects and objects from which one can infer (consistently with the norms of research practices) a disposition to fulfill the targeted use. In other words, a model of a practice includes those properties relevant for defining types for which that practice has a use. Models of practices are ways of organizing whatever information about subjects and objects is relevant to type definition. But the uses themselves are uses that we have reason to care about, because such uses are important to

what various types of subjects do, or at least have the potential to be important to what various types of subjects do.

Now, in the terminology I use, some set of properties sufficient for type definition counts as a criterion of identity for an object or a subject. Models of practices are then criteria of subject and object identity. But there are two important questions here: how to assess how good one's criteria of subject and object identity are for a given practice, and also just what such criteria of identity are criteria of identity for

The model of practices is basically a way of organizing information about uses, which are essential to the criterion of type definition. But the uses themselves are uses that we have reason to care about, because such uses are important to what various types of subjects do, or at least have the potential to be important to what various types of subjects do.

3. A Schema for Assessing the Relevance of Identity Criteria

Given a specific inferential goal, the ontologist can offer reasons for claims about which models of contexts are better or worse. Such reasons would turn on which models of practices are better able to fulfill the inferential goal. If the inferential goal involves, for example, prediction or explanation, the ontologist can assess how well particular models predict or explain the way the subjects of certain practices do in fact behave. Such assessment can in turn make use of norms such as consistency, predictive power, relevance to observable behavior, etc.

Consider, for example, the ontologist who wants to model subjects engaged in shopping in order to figure out what factors motivate those shoppers. The ontologist could adopt

categories of shopping motivation used by consumer researchers. But the ontologist might borrow concepts which anthropologists would typically use in studying shoppers, either in addition to or instead of the concepts used by consumer researchers. The ontologist might turn to anthropologists' concepts in order, for example, to study how shopping behavior contributes to subjects' perceptions of their social class, which might be of less interest to consumer researchers.

Accordingly, the selection of an inferential goal allows the ontologist to choose between various contexts, depending on which contexts are likely to better fulfill that inferential goal. Such an approach is inter-contextual, because the inferential goal defines an expanded context against which one can evaluate the effectiveness of multiple other contexts—say, anthropology versus consumer research—in fulfilling that inferential goal.

Now such a reflective attitude can involve taking the standpoint of a practice or set of practices in order to describe, explain, or predict the behavior of subjects within either the same practice or some other practice or set of practices. The reflective attitude is characteristic of research practices. That is, the ontologist may choose to fulfill an inferential goal by *participating* in practices, like anthropology and consumer research, whose practitioners are themselves concerned with such inferential goals. Practices typically concerned with inferential goals would include, but would certainly not be limited to, academic disciplines and types of research. The more a practice is concerned with an inferential goal, the more "reflective" that practice is. E.g. consumer research is more reflective than shopping. That

does not mean, however, that shoppers might not try to explain or predict their own behavior. Shoppers, at least in their capacity as shoppers, would merely be less concerned with such explanation and prediction than are the consumer researchers whose job it is to study shoppers. For the sake of brevity, one can refer to practices concerned with inferential goals as "reflective practices," even though reflectiveness could more properly be thought of as a spectrum rather than a binary choice between the reflective and the unreflective.

The ontologist can assume the guises of many such reflective practitioners. The guises assumed by the ontologist depend on whichever reflective practices help the ontologist better fulfill the inferential goals with which the ontologist is concerned. Indeed, to achieve a realist agenda, a reflective attitude should dovetail with inter-contextuality. What matters is how well the inferential goal is fulfilled, given realist norms for assessing that fulfillment. The selection of reflective practices as contexts is merely instrumental for this realist agenda.

To offer an account of such an inter-contextual framework for applied ontology, I make use of general concepts from ontology, concepts not dependent on any one type of context. My framework is divided into four step-by-step procedures given in the following three sections.

To better explain, predict, argue about, or describe "what there is," the ontologist would want to accurately and consistently take account of the observable behavior. In the context of practices, such observable behavior involves subjects who have uses for objects. Now the uses which subjects have for objects vary between practices, even when subjects are

using objects of the same type. A biologist may have different uses for salmon than a salmon farmer. The realist ontologist should acknowledge that different types of subjects make different uses of objects of the same type.

The upshot of such variation in uses is that different types of subjects may have reason to draw on different information in identifying the same object type. After all, the information one needs to make some use of an object can depend both on the object type and the use. The biologist mapping the genetics of salmon has at least some different information needs than the salmon farmer trying to make salmon safe and edible. Some empirical research supports this reasoning, and realists ought to care a great deal about empirical research. A salmon farmer, for example, identifies tokens of the type "salmon" in terms of information which the ontologist could model through categories of texture and ease of handling. A biologist, on the other hand, identifies the type "salmon" in terms of categories of taxonomy, phylogeny, and genetics (See e.g. Law and Lien 2012). Indeed, defaulting to the most seemingly "reductive" criterion of identity—e.g. type categories of genetics, phylogeny, and taxonomy instead of texture and ease of handling—may do a poor, unrealistic job of modeling certain practices. I offer extended examples later in this section.

For the ontologist modeling practices, the question is how to identify object and subject types *on a model of a practice*. The realist ontologist's challenge is to ensure that such models accurately and consistently take account of observable differences in uses and information needs. The goal is for such models to do a good job of fulfilling some inferential goal selected by the ontologist. As I discussed in the previous section, such models regiment the information needs of subjects within practices, through categories of property, object, and subject types.

Accordingly, I take on the problem of how to figure out, for a given object or subject type, which criteria of identity are better than others for a given context. What I want to argue is that an inter-contextual approach promotes realist norms. How does the ontologist expand the scope of the context of evaluation, in such a way that the ontologist can evaluate multiple identity criteria and models of practices relative to a broader context? How can that broader context in turn be evaluated according to realist norms? I touched on such questions in previous sections. Now I want to analyze such questions into a series of questions which I explore in the following sections.

Specifically, there are four interrelated questions for the ontologist as a modeler of practices:

(S4.1) How good is a proposed identity criterion for object types relative to the most suitable models of a practice one might select?

(S4.2) How good is such a proposed identity criterion for subject types?

(S4.3) How well has one *selected a practice* for an inferential goal—is the practice one has selected suitable for one's inferential goal?

(S4.4) How well has one *constructed the model* of a practice—how well has one selected the type categories on that model of a practice for one's inferential goal?

Note that how well one answers (S4.1) and (S4.2) can depend on how well one answers (S4.3) and (S4.4). If one has not selected suitable models of practices for assessing a proposed identity criterion, then one may not get a very good answer for how good that identity criterion

is. Recall that subject and object type identity criteria are parts of a model of a practice—a model of a practice is constituted by some finite set of such identity criteria.

The ontologist would likely benefit from common procedures for answering questions (S4.1) through (S4.4). After all, if our ontologist is a realist, she would want to be able to offer good reasons for her choice of context. Indeed, realists are realists in virtue of following the same basic norms such as accuracy, consistency and relevance to observable behavior. Such norms aim to better uncover the truth about some aspect of the world. I hope to show that a common procedure could help the ontologist follow such realist norms across practices. So what kind of procedure would do a good job of fulfilling such norms?

My focus in this section is on question (S4.1)—how to choose better criteria of object identity relative to a given model of a practice. I offer a procedure for answering that question in this section. The following sections discuss how to answer (S4.2) through (S4.4).

To begin, let's have the ontologist take a look at the salmon farmer. The salmon farmer—the type of subject who participates in the practice of salmon farming—has information needs different than those of the biologist (Law and Lien 2012). As a salmon farmer, one's use for salmon is to make those salmon fit for eventual human consumption. According to the first procedure, the ontologist can model the information needed to fulfill a use in terms of a set of categories of types of properties, objects, processes, etc. Together such categories would constitute a *good* criterion of type identity for salmon for the practice of

salmon fishing. So which types of categories *should* model the relevant object type—the salmon—for the practice of salmon fishing?

More particularly, two types of use matter for modeling practices: behavioral uses and inferential goals. A behavioral use is a use that subjects have in physically interacting with object tokens of a given type. Buying, for example, is a certain type of *behavioral* use that shoppers could make of a type of product-object. Stealing or breaking the products would be other such behavioral uses, ones less amenable to the consumer researcher's *inferential goal*. That inferential goal would likely involve figuring out what contributes to consumers buying more products. There may not be a sharp distinction between behavioral and inferential goals, but there is at least a spectrum. The more a given use involves description, explanation, argument, explanation, or some other kind of conceptual work, the more that use warrants being referred to as a "inferential goal."

Accordingly, use determines the scope of the practice that is to be modeled. The aspects of practices which I am interested in modeling here are object and subject type identities. So I begin the first procedure by reducing (S4.1) to the following question (S4.1*): For models of object and subject type identity, how good is some criterion of identity C of some object type O given some type of use U?

For the ontologist modeling practices, the question is how to identify object and subject types *on a model*. The procedure is not asking how to tell which objects count as tokens of the type "salmon." Rather, criteria of *type* identity concern the categories of information needed

to model the identity of a subject or object type for some inferential goal within some specified context. The ontologist would aim to explain, predict, describe, or argue about the categories of information needed to make some use of objects of some type, e.g. “salmon” or “river rapid” or “consumer product.” Such uses would in turn delimit the context—namely the practice—with which the model of subject and object identity is concerned.

The ontologist can thus answer (S4.1*) through the follow procedure:

Procedure 1: Selecting Criteria of Object Type Identity through Practices

P1.1) What information does one need to have about object type O to make use U of objects of type O?

To answer question (1) for U, I argue that the ontologist should offer a model intended to address the following questions:

P1.2) What practice P is most concerned with doing U with O?

P1.3) Which types of ontological categories identify object types on the model of P?

P1.4) To what extent can the proposed type criterion C be modeled in terms of the relations between the categories in 3)?

Let’s say that the answer to question (4) of this procedure (P1.4) measures the *relevance* of identity criteria for subjects or objects relative to some particular use. That is, the

greater the extent to which the criterion can be modeled in terms of the categories in P1.3, the more relevant that criterion is.[11]

Note that the first procedure is intended only to answer (S4.1) by way of (S4.1*). In particular, a separate procedure, which I provide in the following section, would be needed to suggest how the ontologist could do a good job of answering questions (2) and (3) of the first procedure (P1.2 and P1.3). How well the ontologist answers P1.2 and P1.3 depends on how well the ontologist has selected a practice for her inferential goal (question S4.3).

Consistently with the realist project, the model's type categories regiment information about the external world, so as to better fulfill some use the ontologist has in modeling some aspect of that world. Indeed, if the information needs for identifying object types were not *about* object types with real tokens, then the framework would not be very useful to a realist.

Admittedly, one might criticize my use of identity criteria as making "what there is" too contingent on whatever information the ontologist selects. Heil (2002), for example, claims (albeit in a footnote) that identity criteria should merely determine what counts as some sort of

[11] Note also that the conception of effectiveness I introduced in the first section is one which takes account of how well realist norms have been fulfilled. Such norms include relevance to observable behavior, consistency, explanatory potential, etc. The goal of questions (S4.1)-(S4.4), and of the corresponding procedures, is to fulfill such realist norms. If the ontologist has in fact done a good job of answering each question in the procedure (P1.1-P1.4), then my claim is that such norms will have be well-fulfilled in modeling object identity for practices. If object identity has been well-modeled in terms of realist norms, then the resulting model is likely to be of greater use for the ontologist. Indeed, the ontologist would want to use such a model to explain, predict, argue about, describe, or justify certain categories of information as being needed to make some specific use of objects of a given type.

thing, rather than serving as "techniques for identifying." But there is no good reason to say that identity criteria must only be about determining what counts as what, while somehow controlling for and holding constant our information needs in identifying something.

As Horsten (2010) points out, the degree to which identity criteria should serve as a "way of telling" (in Anscombe's phrase) depends on the things we are identifying and our purposes in identifying them. In the context of applied ontology, it does make sense for such identity criteria to conform to our information needs. After all, applied ontology deals with types of things which practices can study and subjects can use. To make use of such entities, however, subjects need information. Indeed, the framework I provide does not allow the ontologist to select whatever information she wishes in modeling an object type. My framework imposes a normative commitment for both selecting and modeling information. The applied ontologist should follow realist norms in selecting that information needed to make some use of an object, while modeling that information in terms of relations between categories of types.

I intend the framework to be used pragmatically, or, if one prefers, "scientifically." That is, the ontologist can test and re-test the procedures for different uses, inferential goals, and practices. The inter-contextual nature of the framework is intended to promote such a scientific approach, by giving the ontologist reasons to switch between models of different practices and meta-practices depending on what yields the best results for the selected inferential goal. The goal is to be able to provide reasons for making such switches. Such reasons could stem from justified or well-reasoned beliefs about which type criteria would yield better results for the inferential goal.

4. Choosing Better Criteria of Subject Identity

To illustrate the selection of categories of subject identity, consider a classic study on shopper types by Westbrook and Black (1985). The inferential goal of consumer research is most generally to figure out what contributes to shoppers buying products. (More particularly, shopper ontologies may allow retailers to better "target" certain products and advertising towards certain types of shoppers [Westbrook and Black 1985, 78]). For that inferential goal, buying products is the relevant behavioral use for shopper subjects; selling products is the relevant behavioral use for retailer subjects. To achieve that inferential goal, consumer researchers have for nearly half a century been occupied with the goal of how to categorize shoppers into types. Consumer researchers would aim to select the types of categories which would be best able to fulfill the inferential goal.

Towards the end of figuring out what gets shoppers to buy products, the categories hypothesized by consumer researchers have included various factors contributing to shopping motivation or shopping behavior (Westbrook and Black 1985). In Westbrook and Black's model, such categories include: "anticipated utility, role enactment, negotiation, choice optimization, power and authority, and stimulation" (1985, 27). From such categories, shoppers were categorized into 6 types. Each category of shopping motivation was related to each type of shopper in terms of the degree of satisfaction which that shopper type derived from that category.

Accordingly, the categories of shopping motivation are included on the model because the consumer researcher has reason to believe that such categories dispose shoppers towards various shopping behaviors. The inferential goal of consumer research would be concerned with explaining or predicting or describing shopping behaviors in terms of such dispositions. The result, on my framework, would be a criterion of type identity specific to each type of shopper. That is, each type of shopper would be identified in terms of a certain set of shopping motivations disposing shoppers to shop in particular ways. One type of shopper, for example, derived high satisfaction from all such categories, such that that shopper type could be characterized as deriving enjoyment from the process of shopping itself (1985, 95). Another type of shopper, the "economic shopper," was more concerned with choice optimization and role enactment—being able to get better deals (1985, 95).

Now criteria of identity which identified shopper types in terms of how long shoppers spent in the restroom, or whether each shopper stood still on the escalator, would likely be less useful for consumer research. That is, a good model of shopping for consumer research would likely include types of categories such as those posited by Westbrook and Black (1985), so as to better achieve the inferential goal of figuring out how to get consumers to buy more products. Accordingly, criteria of identity based only on categories of restroom time and escalator behavior would be assessable as bad relative to Westbrook and Black's model. Even so, the question of how good identity criteria are relative to some model of a practice (S4.1 and S4.2) is different than the question of how well the ontologist has in fact modeled that practice. Is a

given model of a practice really a good model? The latter question (S4.4) is addressed in the next section.

Accordingly, question (S4.2) can be reduced to the following (S4.2*): How good is a given identity criterion C of some *type of* subject S for some *type of* inferential goal M(U) of S? S would, in turn, make some type of use of objects on some practice. Accordingly, on the framework, subject types are participants in practices who have uses for objects. Such subject types would typically be studied by social sciences such as anthropology or economics, though this is not necessarily the case.

Procedure 2: Selecting Criteria of Subject Type Identity through Practices

P2.1) What information does the ontologist need to have about subject type S to make some inferential goal I(U) of S?

To answer (1) for I(U), I then argue that the ontologist should one ask:

P2.2) Which type of use U is S making of objects? (Note that U can be either a behavioral or an inferential goal, but M(U) must be an inferential goal.)

P2.3) Which practice P is most concerned with U?

P2.4) Which subject type categories are associated with U on the model of P?

P2.5) To what extent can the proposed type criterion C be modeled in terms of the categories in (5)?

5. Selecting Practices and Assessing Models of Practices: Two Last Procedures

So far I have suggested that inferential goals are critical for doing a better job of modeling practices, but I have not explained exactly how inferential goals play their role. My basic claim has been that inferential goals allow the ontologist to determine how good her selection of a practice is for examining the subject and object identities in which she is interested. Relative to such an inferential goal, the ontologist can switch between better practices and worse ones. The question for the ontologist is how to tell which practices are better and which are worse for a given inferential goal. Note that question (S1.2) in the first section raised the general question of how to select better and worse contexts.

Accordingly, in this section I aim to address the role of inferential goals. In particular, I offer a procedure for answering question (S4.3), which is a specification of question (S1.2). That is, I offer a procedure for selecting better practices for a given inferential goal. Indeed, in the procedures concerning subject and object identity, I have left open the question of how the ontologist determines which practice is "most concerned" with a given use.

In considering how to better select practices, let's return to the conception of the reflective attitude which I introduced in section 3. Reflective practices—those practices which tend to have more of a reflective attitude—are those practices whose uses for objects and subjects are skewed towards conceptual work such as description, explanation, argument, justification, or prediction. The latter list is the mantra of inferential goal. Academic

disciplines are a paradigm case of reflective practices, though such disciplines by no means exhaust the inventory of reflective practices. By drawing on the concepts of such reflective practices, and perhaps working in the guise of a practitioner such as a consumer researcher or anthropologist, the ontologist aims to devise subject and object identity criteria for practices in fulfillment of an inferential goal. Reflected-upon practices are those practices for which reflective practices have inferential goals. Reflective practices can also take themselves as reflected-upon, as when anthropologists study themselves as practitioners of anthropology.

Towards the end of selecting better practices for an inferential goal, the ontologist must first be able to identify the scope of practices. If the model of a practice is a certain finite set of categories, then the question arises: how are these categories to be selected? How such categories are selected, I have claimed, depends on the scope of the practice. The scope of the practice depends in turn on uses. The practice of shopping is identified through the use which shoppers have for objects—i.e. buying products. Both consumers researchers and anthropologists would likely associate "buying objects" with the practice referred to as "shopping." For those practices which tend to be less reflective, such as shopping and sporting, the uses which define these practices tend to be behavioral, having more to do with physical interaction than with explanation, description, argument, justification, or prediction.

But there is still some ambiguity concerning what is meant by the "scope" of a practice. Consider question P1.2 of the first procedure: For some type of use U, which practice P is most concerned with doing U with object types? Reflective practices would tend to associate certain

practices with certain uses, for the purpose of better studying those practices in accordance with realist norms of accuracy, consistency, predictive power, relevance to observable behavior, etc.

Say that a typical inferential goal of consumer researchers is to figure out what contributes to shoppers buying more products, and a typical inferential goal of anthropologists studying shoppers is to explain how shopping contributes to the formation of social class identity. Both anthropologists and consumer researchers trying to offer an accurate and consistent account of shopping would likely associate the practice of shopping with the behavioral use of buying products. Where such reflective practices as anthropology and consumer research differ is in their inferential goals, and in the categories of subject and object identity selected for the reflected-upon practice in order to best fulfill such inferential goals.

So multiple reflective practices, in following realist norms, may associate a given behavioral use with the same reflected-upon practice. But that is not to say that such reflective practices would associate the same *set* of behavioral uses with the practice they reflect upon. It is not even to say that some reflective practice might not differ from many other reflective practices in *not* associating some behavioral use with the practice that is reflected upon. Anthropologists might associate with shopping the behavioral uses of buying products, socializing with other people, and acting out certain social class roles. Consumer researchers might focus only on buying products, or might at least privilege buying products in a way that anthropologists would not, e.g. by defining the inferential goal of consumer research relative to

buying products. Moreover, the ontologist might devise some inferential goal for shopping which could be accurately and consistently carried out *without* even focusing on the behavioral use of buying products. Perhaps describing or explaining or predicting the behavior of shoppers might, for that different inferential goal, not even involve shoppers buying products.

Indeed, perhaps the ontologist has an inferential goal concerned only with making the practice of shopping as secure as possible. For our ontologist qua security researcher, the inferential goal is not figuring out what contributes to shoppers buying products, but what contributes to shoppers stealing products, and how such thefts might be reduced. The ontologist might then categorize shoppers in terms of dispositions to steal, rather than dispositions to buy.

Accordingly, I suggest that the ontologist can leave the boundaries of a practice indeterminate, insofar as the ontologist does not need to have a finite set of all the uses with which a practice might be concerned. If the essence of a practice is given by such a finite set of uses, then one need not argue that practices have essences in order to answer question P1.2. An ontologist need only have a finite list of the categories of subject and object identity for a given model of a practice. But the construction of such models does not entail that the models are meant to explain all possible uses which could plausibly be associated with that practice.

So I maintain that the ontologist should ask the following set of questions in order to answer step 2 of the first procedure (P1.2) and (P2.3): Which practice is most concerned with making some use or set of uses U of object types?

Procedure 3: Selecting the Practice

P3.1) What inferential goal M(U) does the ontologist wish to make of subjects doing U?

P3.2) Can the ontologist identify some reflective practice or practices concerned with doing M(U)?

P3.3) Is there some common practice by which individuals engaging in U would be described by the reflective practices in (P3.2)?

Procedure 4: Modeling the Practice

If the answer to (P3.3) is yes, then the ontologist should ask:

P4.1) Which categories of subject and object identity do the reflective practices in (P3.2) associate with the practice in (P3.3)?

P4.2) Which are the criteria best justified, according to realist norms, for assessing how well M(U) has been fulfilled?

P4.3) Do the categories in (P4.1) do a good job of fulfilling M(U) according to the criteria of assessment in (P4.2)?

If the answer to either (P4.3) or (P3.3) is no, then the ontologist should ask:

P4.5) Are there any categories of subject or object identity used by any of the reflective practices in (P3.2) which the ontologist has reason to think would do a good job of fulfilling M(U) according to the criteria of assessment in (P4.2)?

P4.6) Are there any combinations of the categories in (P4.5) which do in fact do a good job of fulfilling M(U) according to the criteria of assessment in (P4.2)?

If the answer to (P4.6) is still no, then the ontologist may have to resort to fashioning new categories of subject and object identity. Fashioning such new categories would seem like a mysterious process, but I believe it is at least normatively guided. With the end in view of fulfilling an inferential goal according to certain realist norms, the ontologist would have to experiment with linking together various concepts into such new categories of subject and object identity. Such new categories would then have to be tested again for how well they fulfilled the inferential goal according to realist norms.

In the end, this fashioning of new categories would likely involve modeling a practice which would previously have been overlooked by contemporary reflective practices. In the process, a new reflective practice could emerge. Indeed, one can take the history of philosophy itself as a paradigm case of the emergence of new reflective practices concerned with modeling previously overlooked practices. One need only look to physics, psychology, economics, and even, perhaps, applied ontology to find examples of such reflective practices.

Conclusion: The Importance for Ontology of Taking Account of What Research Questions are About

In sum, if ontology is to better specify the results of empirical research, ontologists need a way to take into account what that research is about. My claim may seem very simple, and I think it is. But insofar as Quinean ontology lacks a criterion for selecting better theories and categories of information for a inferential goal, Quinean ontology does not offer good reasons to select different theories and categories of information based on what a given research question is about.

In the cases I have discussed, the Quinean focus on scientific theory does not yield the information which best fulfills the research question or inferential goal. The problem is that without a criterion for selecting better theories or categories of information, the ontologist would not be able to *defend* her reasons for selecting some categories of information about an object type as better than other categories relative to a given behavior involving that object type. Indeed, different research questions require different categories of information in order to be effectively fulfilled, especially in regard to differences between the social and natural sciences. Returning yet again to Law and Lien's study (2012), the researcher seeking to explain or describe salmon, for the behavioral context of farming salmon, would want to draw on categories of texture and color and vaccination status, rather than categories of phylogeny and taxonomy.

I have accordingly argued that ontology should take account of empirical research into a variety of social practices, not only the natural sciences. Taking account of behavioral context more fully commits ontology to specifying the results of empirical research across all the possible domains of social practices. But to account for such a variety of behavioral contexts, with so many possible categories of information, ontologists should have a criterion for choosing better categories of information.

Most of the categories of information I have discussed include dispositional properties of object types. Indeed, the research questions I have considered focus on those dispositions which make object types useable for the behavior being studied, such as the dispositions which make salmon farmable.

To defend some categories of information about object types as doing a good job of fulfilling a research question, the ontologist should have a criterion that considers object types to be *well*-identified through dispositions which effectively explain the behavior being researched. Such a criterion may be quite simple. Yet the Quinean focus on scientific theory has obscured the importance of there being such a criterion in the first place.

The lack of a criterion for choosing better categories of information may not be such a big problem for research questions in which the ontologist only needs to draw on the best scientific theory available. But there are likely to be many research questions, particularly those involving social practices other than the sciences, where drawing on scientific theory provides information that fails to do a good job of answering the research question. One

possible solution is to say that ontology need only be concerned with issues of scientific theory. But I have argued that there is interesting research into social practices other than science with which ontologists also have good reason to be concerned. There is thus also good reason for ontology to have a criterion for better identifying object types in terms of what different research questions are asking about.

Bibliography

Asher, Nicholas. *Reference to abstract objects in discourse*. Vol. 50. Springer Science & Business Media, 2012.

Bird, Alexander. *Nature's metaphysics: Laws and properties*. Oxford University Press, 2007.

Bird, Alexander. "Natural kinds." *THE ROUTLEDGE COMPANION TO METAPHYSICS* (2008): 497.

Boyd, Richard (1999). Homeostasis, species, and higher taxa. In R. A. Wilson (ed.), *Species: New Interdisciplinary Essays*. MIT Press. pp. 141-85.

Brandom, Robert. *Making it explicit: Reasoning, representing, and discursive commitment*. Harvard university press, 1994.

Gordon, I. Herbert. “Classification of Rapids, Water Level and Canoeists.” *Paddling.com*. <https://paddling.com/learn/classification-of-rapids-water-level-and-canoeists/>. Accessed 14 December 2017.

Haugeland, John. "Truth and finitude: Heidegger’s transcendental existentialism." *Heidegger, authenticity, and modernity: Essays in honor of Hubert L. Dreyfus* 1: 43-77. MIT Press, 2000.

Hawley, Katherine, and Alexander Bird. "What are natural kinds?." *Philosophical Perspectives* 25.1 (2011): 205-221.

Heidegger, Martin. *Being and time*, trans. Joan Stambaugh. State University of New York Press, Albany (1996).

Heil, John. *From an ontological point of view*. Oxford: Oxford University Press, 2003.

Horsten, Leon. "Impredicative identity criteria." *Philosophy and Phenomenological Research* 80.2 (2010): 411-439.

Kohn, Eduardo. "Anthropology of ontologies." *Annual Review of Anthropology* 44 (2015): 311-327.

Kuhn, Thomas. "Objectivity, value judgment, and theory choice." *Arguing About Science* (1977): 74-86.

Kuhn, Thomas. "Logic of discovery or psychology of research." *Criticism and the Growth of Knowledge* (1970): 1-23.

Kuhn, Thomas. *The Structure of Scientific Revolutions*. University of Chicago Press, 1996.

Law, John, and Marianne Elisabeth Lien. "Slippery: Field notes in empirical ontology." *Social Studies of Science* 43.3 (2013): 363-378.

Merleau-Ponty, Maurice. Phenomenology of perception. Routledge, 2013.

Munn, Katherine, and Barry Smith, eds. *Applied ontology: an introduction*. Walter de Gruyter, 2009.

Palecek, Martin, and Mark Risjord. "Relativism and the ontological turn within anthropology." *Philosophy of the Social Sciences* 43.1 (2013): 3-23.

Peirce, Charles Sanders. *Collected papers of charles sanders peirce*. Vol. 2. Harvard University Press, 1974.

Quine, W. V. "On empirically equivalent systems of the world." Erkenntnis 9.3 (1975): 313-328.

Quine, W. V. "On what there is." *The Review of Metaphysics* 2.1 (1948): 21-38.

Quine, W. V. "Ontological relativity." *The Journal of Philosophy* 65.7 (1968): 185-212.

Quine, W. V. "Relativism and absolutism." *The Monist* 67.3 (1984): 293-296.

Quine, W. V. "Speaking of objects." *Proceedings and addresses of the American Philosophical Association*. Vol. 31. American Philosophical Association, 1957.

Reckwitz, Andreas. "Toward a theory of social practices: a development in culturalist theorizing." *European journal of social theory* 5.2 (2002): 243-263.

Sartre, Jean-Paul. Being and Nothingness: A Phenomenology Essay on Ontology. Washington Square Press, 1992.

Schaffer, Jonathan. "Truthmaker commitments." *Philosophical Studies* 141.1 (2008): 7-19.

Searle, John R. "Social ontology: Some basic principles." *Anthropological theory* 6.1 (2006): 12-29.

Shove, Elizabeth, Mika Pantzar, and Matt Watson. *The dynamics of social practice: Everyday life and how it changes*. Sage, 2012.

Smith, Barry. "Ontology." In Luciano Floridi (ed.), *Blackwell Guide to the Philosophy of Computing and Information*. Oxford, Blackwell: 155-166 (2003).

Smith, Barry and Bert Klagges. “Bioinformatics and Philosophy.” In *Applied ontology: an introduction*. Walter de Gruyter, 2008: 21-39.

Van Inwagen, Peter. "Meta-ontology." *Erkenntnis* 48.2 (1998): 233-50.

Wittgenstein, Ludwig. *Philosophical investigations*. John Wiley & Sons, 2009 [orig. 1953].

Wetzel, Linda. *Types and tokens*. MIT Press, 2009.