

Normativity and Experimental Philosophy

Shannon S. Abelson

Thesis submitted to the faculty of the Virginia Polytechnic Institute and State University in
partial fulfillment of the requirements for the degree of

Master of Arts
In
Philosophy

Joseph C. Pitt
James Klagge
Kelly Trogdon

March 24th, 2015
Blacksburg, Virginia

Keywords: experimental philosophy, normativity, practical knowledge

Normativity and Experimental Philosophy

Shannon S. Abelson

ABSTRACT:

The normative conception of philosophical methodology take as a starting point the teleological prioritization of practical knowledge over theoretical knowledge I examine what a normative philosopher might make of the emergent practice of experimental philosophy. Generally speaking, experimental philosophers set as their methodology the suspension of a priori reflection in favor of empirical experimentation in order to examine the cognitive factors that influence concept application and behavior. I begin by examining the normative view of philosophy. This view is rooted in a pursuit of practical knowledge, which is an inherently normative endeavor. I then consider how a normative philosopher might approach the role of X-phi. I examine two prominent X-phi projects in order to establish some criteria for meeting that normative philosophical expectation. I conclude that some X-phi projects may not meet the normative expectation because of a pervasive neglect of practical knowledge, while others do manage to incorporate a practical element and fall within the bounds of the normative conception.

Acknowledgements

I would like to express my sincere gratitude to the members of my committee, Joseph Pitt, James Klagge, and Kelly Trogdon for reading and offering helpful feedback on this project and for their continued support. I would also like to thank the members of my cohort, especially Steven Mischler and Caitlin Parker reading earlier drafts and speaking with me about this project. I also wish to thank my mother, Michelle Maggs, for her enduring support for my academic development.

Table Of Contents:

Introduction	1
Section 1: The Normative Conception of Philosophy	4
Section 2: What <i>is</i> X-phi and What Ought One do with it?	10
Section 3: Two Experimental Approaches	15
Section 4: Philosophical Usefulness	22
Conclusion	26
References	28

Introduction

Philosophers who embrace what has been called the “normative turn”¹ in philosophical methodology take as their starting point the teleological prioritization of practical knowledge over theoretical knowledge.² This normative conception of philosophical practice influences how they may perceive the role of certain philosophical projects. I wonder what a normative philosopher might make of the emergent practice of experimental philosophy. Generally speaking, experimental philosophers set as their methodology the suspension of a priori reflection in favor of empirical experimentation in order to examine the cognitive factors that influence concept application and behavior.³ Many practitioners of experimental philosophy argue that X-phi, as it has been nicknamed, is a more traditionally faithful way of doing philosophy that hearkens back to the historical question of how the mind works, while simultaneously espousing it as a novel movement. I begin by examining the normative view of philosophy. This view is rooted in a pursuit of practical knowledge, which is an inherently normative endeavor: the pursuit of navigation through the “logical space of reasons.”⁴ I then consider how a normative philosopher might approach the role of X-phi. Recent attempts to address X-phi’s role in the philosophical landscape have tried to configure it as a type of philosophical naturalism. I propose that this characterization of X-phi does not account for an important sense in which some X-phi projects fall short of

¹ This term has been used by Robert Brandom and James O’ Shea.

² I would like to thank Joseph Pitt, James Klagge, Kelly Trogdon, Lydia Patton, Tristram McPherson, Steven Mischler, and Caitlin Parker for their helpful commentary on this project.

³ This is meant as a general, non-exhaustive claim. By “concept application,” it is generally meant the situations in which a competent user of a language decides to employ certain concepts. It is not necessary here to proscribe any particular theory of meaning. I choose the term “cognitive factors” in order to remain neutral on the issue of the philosophical significance of intuitions.

⁴ This is the conceptual apparatus by which we submit and evaluate the reasons that underlie our knowledge of the world and ourselves in it.

a normative philosophical expectation. This expectation involves a teleological prioritization of practical knowledge. I consider two prominent X-phi projects in order to establish some criteria for meeting that normative philosophical expectation. I conclude that some X-phi projects may not meet the normative expectation because of a pervasive neglect of practical knowledge, while others do manage to incorporate a practical element and fall within the bounds of the normative conception. The former projects, while failing the normative expectation, may still meet a standard of philosophical usefulness. The exact nature of this usefulness will hinge on the type of contribution X-phi projects make to philosophical discussions. If one assumes a normative conception of philosophy, X-phi's discovery of knowledge about the sources of people's intuitions, behavior, language, and neurological states does in fact contribute significantly towards populating the landscape of the logical space of reasons, but what the normative philosopher ultimately seeks is a compass.

In what follows I appraise and situate X-phi's place in the philosophical landscape according to a conception of philosophy as an inherently normative enterprise. I will not in this paper argue for the acceptance of this view. My project is conditional in nature. If one were to accept the normative conception of philosophy—and I think there are powerful reasons why one should—then certain things follow regarding how one should view experimental philosophy. I proceed under the assumption that philosophers ought to seek practical knowledge foremost, and that practical knowledge is attained by a fusion of knowing-how and knowing-that, where knowing-how takes primacy. In other words, the normative philosophical orientation requires that knowledge-that or theoretical knowledge, be sought in service to the search for practical knowledge or

knowledge-how. Practical knowledge should be prioritized. If an X-phi project is to count as being normatively oriented in the appropriate sense, it must ultimately seek practical knowledge.

Section 1: The Normative Conception of Philosophy

Under a specific, though by no means peripheral interpretation of the aim and methodology of philosophy, philosophy is considered a principally practical endeavor. This is the conception of philosophy that has been explicitly advocated by Wilfrid Sellars, though it can be seen in the work of others.⁵ It is useful here to take as a starting point the now famous and oft-cited quote from Sellars: “the aim of philosophy, abstractly formulated, is to understand how things in the broadest possible sense of the term hang together in the broadest possible sense of the term” (1962, 7). Sellars has in mind here a range of “things,” some of which may be naturalistically explained, while others are not so easily integrated into a naturalistic picture. The goal of the Sellarsian is to take stock of the philosophical landscape and attempt the construction of a procedural map, by which she may navigate the various connections between seemingly disparate subjects and subject matters. The contents of the philosophical landscape include not only those imperceptible entities postulated by our best scientific theories, but also those whose existence is presupposed in our common-sense conceptual apparatus. Roughly speaking, for Sellars the philosophical landscape includes all those things, material or immaterial, about which we may speak.

According to Sellars, upon reflection we are confronted with two conflicting pictures of the world. There is the picture that comes most intuitively, that of the world as it directly relates to us. This manifest image of man in the world is a conceptual framework in which human beings are aware of themselves as being human beings in the world. What this amounts to is complex, but it essentially includes all perceptible

⁵ Normative themes of this type are also present in the work of Friedrich Nietzsche and Ludwig Wittgenstein’s later work, among many others.

artifacts of humanity's intellectual history and development. Under the manifest image, the ocean is blue, the Vatican is located inside Rome, and persons have intentions, emotions, thoughts, and conversations. Conversely, the scientific image (itself a conglomeration of scientific views of the world, such as those theories postulated in theoretical physics, chemistry, etc.) postulates imperceptible entities and forces in order to explain the existence and behavior of physical objects. The ocean is not blue in the scientific image; in fact there is no "ocean" at all. There is a swirling mass of imperceptible particles, the arrangement of which is determined by imperceptible forces. The scientific image attributes ontological supremacy to what our best scientific theories quantify over; it thus presupposes a highly reductive form of scientific realism. Those who become reflectively aware of this conflict will recognize the seeming incompatibility of these images, while simultaneously recognizing the necessity of their reconciliation. To attempt to situate one's claims in the context of these two images is to place them in the logical space of reasons (Sellars 1962).

Now I turn to the normative view of philosophy derived from this perspective, Navigating this conflict—learning how to integrate one's claims in the logical space of reasons—requires practical knowledge. Projects that seek this practical knowledge are normative ones. As mentioned earlier, by practical knowledge I mean knowledge-how. A paradigmatic piece of practical knowledge might be a recommendation on how to behave in a certain situation or how to do something (like riding a bike) or a recommendation that someone should believe A under circumstances B (1949, 209). Usually the philosophical method that is associated with practical knowledge is pragmatism. But while what I am examining is undeniably a pragmatic aim, it is intended to be a broad

one. The appeal of the normative conception lies in its promise to reveal a way of navigating the logical space between the manifest and scientific images. It thus proposes to facilitate a certain kind of intellectual action. This conception will exclude information that is idle with respect to facilitating intellectual action.⁶ One way of conceiving of intellectual action may be to imagine what one would require intellectually in order to solve a certain problem.⁷ If I aim to give a formal proof of an axiom, most useful to me will be knowledge of the rules at my disposal. Such knowledge facilitates my action (my giving the proof). Practical knowledge plays a dual role: it both facilitates intellectual action and functions as source of justification for one's acceptance of a position. If asked how I derived an axiom, the appropriate response is to recount the steps I took to prove it which will entail a citation of the rules that enabled me to do so. As such, much of what goes on in metaphysics, philosophy of mind, philosophy of religion, aesthetics, and so forth, that aims to discover how to proceed in these respective domains, will be encompassed by the normative conception.

So I put forward the following contention: normative philosophical orientation requires that a project aim foremost to provide practical knowledge. It is important here to note several things. First, and most importantly, this does not exclude projects that provide propositional or theoretical knowledge from being philosophical. Theoretical knowledge is important, even integral; it's just not the dominant aim in a normatively oriented project. In fact, in a large number of philosophical works, practical knowledge may not be attainable without propositional knowledge. I do not have in mind here that

⁶ I wish to remain neutral with regard to what theory of mental action one should espouse, though the account I am exploring does require commitment to some theory of mental action (eliminativist accounts of the mental are incompatible with the normative conception).

⁷ Though problem-solving does not exhaust the applications of intellectual action.

one must know the mechanics of a bicycle in order to know how to ride it. But for example, when attempting to articulate a prescription for how one should justify a belief, a philosopher should have some working knowledge of possible sources of belief. Again, this is an argument that originates with Sellars. Sellars describes the philosopher as having a responsibility of keeping her “eye on the whole,” rather than being simply a “reflective specialist” (1962, 7-8). In order to have one’s eye on the whole—to know one’s way around—one must know a sizeable amount about a great many subjects: “reflective knowing one’s way around in the scheme of things presupposes a great deal of reflective knowledge of truths (1962, 8).” What is meant by the “whole” is just as ambitious as one might imagine. The “whole” includes any subject and its subject-matter that has any bearing on the world, persons, their composition, behavior, intentions, etc.⁸ Philosophers seek a tenable way of navigating around this whole and so must possess the requisite amount of theoretical knowledge about its constituents. Such knowledge functions as a set of composite pieces of the larger normative project of developing the practical knowledge needed to navigate the philosophical landscape.⁹ It is important to note as well that even projects that provide bad or untenable practical knowledge can count as being philosophical. A work does not cease to be normatively oriented simply because no one thinks its viable; bad philosophy is still philosophy.

One of the difficulties in determining what will and will not count as a piece of practical knowledge under the normative conception is that philosophers often appeal to

⁸ But to say that one must keep her “eye on the whole,” stops short of saying that one should “keep her eye on everything.” Presumably there will be some things that are simply philosophically irrelevant. An examination of what such things are irrelevant is another project.

⁹ I am comfortable asserting, at least tentatively, that the normative conception is also compatible with epistemic views such as that of Robert Brandom, that conceive of all or almost all knowledge as practical. Such views hold that theoretical knowledge is either functionally practical or composed of practically acquired knowledge.

wide variety of evidential sources to support their arguments. It seems like any attempt to give a general account of practical knowledge runs the risk of excluding something potentially important. One way to address this difficulty can be to imagine that what counts as being a piece of practical knowledge, what functions to aid intellectual action, is a kind of tool. So we distinguish between what a philosopher may use (the subject matter of her project) and what actually aids her in a project. There may remain some ambiguous cases, but application of this criterion could potentially mitigate that ambiguity in many of them.

At this point it may be prudent to pause and consider what amounts to a major point of controversy with the normative conception as it has been articulated by Sellars. While I have suggested that this approach has great merit, it remains the case that this approach will be quite exclusive. Insofar as a work does not demonstratively both seek a certain piece or pieces of practical knowledge *and* prioritize that knowledge-how over any propositional knowledge sought, that project will fail to count as a philosophical project. I have certain misgivings about the exclusivity of this conception, as I imagine many others will. My hesitance to embrace this approach in its purest form stems primarily from my concern that many projects may be seen (or are expressed as) being practically-oriented, but do not explicitly distinguish between types of knowledge, practical or propositional. It is fairly clear that Sellars asserts a qualitative distinction between practical and propositional knowledge, but I am less convinced that a difference in type is sustainable. Consider again how one may distinguish between practical and propositional knowledge. A paradigmatic piece of practical knowledge may be a prescription: given the set of conditions X, one should adopt Y, where, for example, X

might be a set of observations about neurological processes and Y may be a belief that mental events are reducible to functional states of the brain. Both X and Y in this case are pieces of propositional knowledge or knowledge-that. They describe a state of affairs. Then the piece of practical knowledge is the prescription derived from these propositional contributions: that should Y adopt conception of mental states.

But does the prescription count as relevantly different in type from the propositional components? It appears one could answer either way. It is different insofar as what it requires from the agent is, as I have put it, some kind of intellectual action. But one could easily imagine that the presentation of some piece of propositional knowledge can require that the agent either assent or dissent. Perhaps the difference consists in one's ability to suspend judgment as to the truth of a propositional statement, but then I can also imagine one suspending agreement with a prescription. This is not an issue that is easily settled. However, in the interest of reformulating the normative conception to be more sensitive to these concerns, I propose that the distinction between practical and propositional knowledge may be replaced by a notion of practical or propositional orientation towards one's project. Under this weak normative conception, a project that emphasizes the practical significance or implications of its contributions is within the bounds of the normative conception. It may turn out to be the case that there is a substantive distinction to be made after all between knowledge-how and knowledge-that. But the normative conception can be expanded to include both projects that include a teleological prioritization of knowledge-how or practical knowledge (strong version) *and* projects that recognize (explicitly or implicitly) and encourage reflection on the practical implications of their contributions (weak version).

Section 2: What is X-phi and What Ought One do with it?

Having set forth the criteria for the normative conception, I now turn to consider the role of experimental philosophy and whether under a recent characterization it satisfies the normative requirement. In a recent interview Neil Levy points out that X-phi projects use such diverse methods as the administration of opinion polls, observation of behavior, and the study of brain states (Levy and Kitano 2011). David Rose and David Danks have defended a similar “broad view,” in which X-phi should be seen as a collocation of philosophical naturalism and cognitive science (2013). Rose and Danks argue that many of the objections to experimental philosophy and characterizations of it as novel movement operate under a narrow definition. Under the narrow view, experimentalists carry out empirical investigations of some sort in order to determine the sources of peoples’ intuitions. Rose and Danks contend that this definition is problematic for numerous reasons and should be set aside in favor of a more defensible and historically-conscious broad view, wherein experimental philosophy is seen as a “co-location in the same body of i) philosophical naturalism; and ii) standard, everyday cognitive science.” (2013, 2)¹⁰ They define philosophical naturalism as a position from which one commits to examine empirical evidence in the course of drawing philosophical conclusions. They go on to argue that the historical presence of philosophical naturalism conflicts with X-phi’s reputation as a novel movement and that

¹⁰ Perhaps worryingly, Rose and Danks do not provide a definition of what they mean by “standard, everyday cognitive science.” They refer interchangeably to psychological experimentation and empirical investigation. They exercise a bit more care in refining what may count as an experimental philosophy project, but defend their “broad view” only from charges of excessive broadness with regard to their definition of philosophical naturalism. There is no definitive articulation of “cognitive science” or qualification of whether the notion of a cognitive science experiment is either too vague or too broad.

it is only a sociological instance of a coupling of philosophical naturalism and the actual conducting of experiments.

It is not clear that Rose and Danks have provided a characterization of philosophical naturalism that would satisfy the normative philosopher like Sellars. While there is widespread disagreement about what philosophical naturalism entails and it is doubtful whether a set of necessary and sufficient conditions for it could be given, there does seem to be a bit more that we need say than “the philosophical commitment that particular types of empirical facts, as determined by one or another science, are relevant to particular philosophical questions or problems” (2013, 8). What makes the commitment philosophical? Very few would argue that X-phi is not naturalistic in a strictly empirical sense. It is uncontroversial that what is sought when conducting empirical investigations is an explanation, maybe contributing to or against the best current scientific theory of the phenomenon in question, the creation of a new theory, exploring the possibilities for further research, and so forth (Franklin 2007). However, to stop here may lead one to a precarious position wherein she is unable to distinguish between a philosophical experiment and one conducted by scientists with no professed philosophical interests in mind. Are all works featuring an experiment classified as experimental philosophy? This issue might not be of particularly concern to many philosophers. For instance Quine, who sees philosophy and science as continuous would have no problem with a definition of philosophical naturalism that blurs these lines (Quine 1969). But a normative philosopher who sees the aim of philosophy as the navigation of the naturalistic (scientific image) and the normative (manifest image) might be troubled by this implication.

The reason why such an implication would be undesirable to a normative philosopher can be captured in two ways, depending on whether one adheres to the strong or weak version of the normative conception. Under the strong version, a project must possess an identifiable and prioritized piece of knowledge-how. As previously mentioned, knowledge-how is meant to play a justificatory as well as enabling role: not only does one need it to navigate the philosophical landscape, it is what she appeals to in order to justify her position. Scientific projects are for the most part geared toward explanation rather than justification. On the weak version, the barrier between explanation and justification may be slightly more forgiving. The distinction between normative and non-normative projects will hinge on the orientation of the author and her treatment of the propositional content. In the case where she makes no attempt to elucidate the justificatory role her findings might play, the project is non-normative. This criterion is elastic enough that some scientific projects may count as normative (perhaps an example of this may be found in theoretical physics, when one tries to reconcile a classical conception of the law of gravity with quantum mechanics). In any event, the definition of philosophical naturalism provided by Rose and Danks specifies neither the strong nor the weak criteria.

In fact, there are a large number of X-phi projects that seem to operate under this non-normative version of philosophical naturalism. Joshua Knobe, in his defense of a conception of X-phi as being cognitive science, claims that the entrance into any experimental project with the presuppositions of a pre-existing theory will unnecessarily bias the experiment. X-phi projects examine the (seemingly isolated) *effects* that may influence concept application (Knobe 2014). Knobe has in mind here the conceptual

frameworks inherited from conceptual analysis. He elsewhere claims that experimental philosophers may instead appeal to the frameworks developed in the tradition of cognitive science, perhaps such as a computational theory of mind. But that does not mean that these frameworks do not themselves contain a conceptual framework (in a computational theory of mind, a conceptual framework will contain stipulations as to what “computation” entails, notions about formal symbol manipulation and so forth). The same potential for bias is present with an appeal to them. So if bias is to be avoided, it seems one would also need to steer clear of the frameworks of cognitive science. A hesitance to engage the theoretical frameworks that may provide a normative background to many of these experiments will make it harder to distinguish X-phi projects from non-philosophical experimentation.

There is seldom identification of such informing theories in these types of experimental projects. They note an effect, they construct a test to determine its presence or absence, and then they *end*. Often, the experiments conclude with recommendations for further research. This may introduce an opportunity for them to fuel other practically oriented projects, perhaps by using their results to inform potentially novel theories of cognition. From those recommendations, another philosopher may develop a theory to account for the effect. But this most often happens externally to the original experimental project. X-phi projects often do not *internally* attempt to proffer new theories. This is a case where the results of an X-phi project prove useful for further philosophical work, but it is presumptuous to imagine that just because the results of an experiment can subsequently lead to a larger theory (or philosophical project), that experiment has the distinction of being theoretical (or philosophical). Under the

normative conception, for Knobe and those experimentalists who follow his lead a dilemma arises: either the experimentalist attempts to invoke as little theoretical or philosophical frameworks as possible and risks blurring the line between an experimental philosophy project and a non-philosophical experiment, or she invokes such frameworks and risks the undesirable bias that may accompany such a move and undermine the objectivity of her experiment.

Section 3: Two Experimental Approaches

For an illustration of this dilemma, consider one of the most famous X-phi experiments, Knobe's "Intentional Action and Side-Effects in Ordinary Language," which carries out two experiments, each designed to determine whether or not the application of the concept "intentional" is correlated with a belief that the subject caused harm rather than helped. Knobe and his associates interviewed philosophical laymen in a public park in a project designed to uncover their intuitions regarding two related thought experiments. In the first instance, 78 participants were asked to consider the following thought experiment:

The vice-president of a company went to the chairman of the board and said, "We are thinking of starting a new program. It will help us increase profits, but it will also harm the environment." The chairman of the board answered, "I don't care at all about helping the environment. I just want to make as much profit as I can. Let's start the new program." They started the new program. Sure enough, the environment was harmed (Knobe 2003).¹¹

They were then presented with an alternative version of the experiment, wherein the only detail changed was that the environment was helped rather than harmed. The participants were then asked to rate on a scale of 0-6 how much blame (in the former scenario) and how much praise (in the latter) they assigned to the chairman. The participants were far more likely to assign blame than praise to the chairman. Knobe takes this as strong prima facie evidence that people correlate intentionality with blameworthiness and not with praiseworthiness. He concludes based on the experiments that there is an asymmetry in

¹¹ An additional experiment with comparable details was subsequently conducted with 42 more participants.

whether people call an act “intentional” in that they are more likely to attribute intentional behavior to acts that cause bad side effects (Knobe 2003).

The results of this experiment may raise interesting questions. They may spur one to rethink whether attributions of intentional behavior can be impartial. They may also, as Knobe suggests, give one cause to look further into whether the side effects of an action influence other ascriptions. These results may be quite interesting to philosophers of mind and language. But it is apparent that there is nothing *internal* to the paper that attempts to situate these results in a larger theory of mind, action, or language. Knobe does situate the experiment in debates about the concept of intentional action (albeit briefly), but he does not take the natural next step; he does not return to the debate after discussing the experiment to speculate as to how his results may figure into the philosophical landscape. He does not introduce a practical element (much less does he let it be the dominant aim).

Certainly this X-phi project and others like it conform to Rose and Danks’ broad view.¹² They acknowledge the importance of empirical data to answer a question about concept application and they involve the actual conducting of an experiment. As might be expected, they do not explicitly commit to the importance of a larger theory or framework behind their results. However, operating under the assumption that philosophical naturalism entails additional reflection on a wider philosophical picture, which includes the provision of knowledge-how under the strong conception, or an

¹² For the purpose of conserving space, I do not examine more projects in depth here. But a similar critique can apply to such papers as: Knobe, Joshua and Erica Roedder. (2009); Nadelhoffer, Thomas (2006); Knobe, Joshua and Ben Fraser (2008); Beebe, James. R., and Sackris, David (2010).

orientation towards the practical implications of the findings under the weak one, the project still does not meet the normative expectation. The project does not include an attempt to bring these results to bear on philosophical questions in any robust way. Whether others use the findings to do so in subsequent papers is tangential; such external applications point to philosophical usefulness but do not guarantee that the projects are themselves philosophical under the normative conception. It does not provide an aid for intellectual action.

The problem I have highlighted is that the incorporation of X-phi into philosophical naturalism is tenuous at best, and at worst requires a relaxation of naturalism to the point that it is both barely philosophical and seems designed to fit the current state of some X-phi projects. Someone who prefers to reject the assumption I have introduced, that philosophical naturalism requires attention to some framework, is free to do so, but they may have to bite the bullet when it comes to distinguishing an experiment conducted by a scientist with no philosophical ambitions and one conducted by an experimental philosopher (a concession the normative philosopher will not be willing to make). Under Rose and Danks' broad view, there is no easy way to distinguish them. Rose and Danks claim to avoid this problem by making the qualification that X-phi projects are naturalistic, which they define as an appeal to empirical data to bear on philosophical questions. But they do not satisfactorily articulate what these philosophical questions may be, what makes a naturalistic commitment philosophical, and, even if they did, they would have to show that a significant number of X-phi projects do this, which I suspect is not the case.

Rather, I suspect that the threat of bias Knobe fears is chimerical. It seems perfectly reasonable to trust that a philosopher who conducts an experiment designed to bear on philosophical questions may appeal to certain frameworks without risking bias in the process. When an ethicist examines the tenability of a moral standard, she is free to explore how it holds up against a Kantian or Utilitarian set of criteria without unnecessarily biasing herself. Indeed we tend to hope her professional training in philosophy and her intellectual rigor ensure that she will not be easily biased, though it is of course possible.

Consider another prominent X-phi project, one that does engage existing theoretical and philosophical frameworks, does not seem to compromise objectivity, and would fit the criteria I have given for the normative conception of philosophy. “Against Arguments from Reference” and its earlier incarnation “Semantics, Cross Cultural Style,” by Edouard Machery, et al. was constructed against the backdrop of pervasive appeal to arguments from reference in analytical philosophy. According to Machery, et al., intuitions factor centrally to theories about reference. They contend that philosophers use the method of cases in order to determine the correct theory of reference. If they are right about this fact, then it follows that if there were substantial discrepancies in intuitions among different groups, a correct theory of reference would be difficult or impossible to attain. Moreover, because so many other debates in philosophy rely on some argument from reference being true, the rejection of arguments from reference could have widespread ramifications (Machery, et al. 2009).¹³

In order to illustrate their case, the experimentalists appealed to a study they conducted in which American and East Asian undergraduates were presented with

¹³ They specifically cite the debates around eliminative materialism and scientific realism.

several cases (variations on Kripke's Gödel case) designed to test their intuitions with respect to reference. For example,

Suppose that John has learned in college that Gödel is the man who proved an important mathematical theorem, called the incompleteness of arithmetic. John is quite good at mathematics and he can give an accurate statement of the incompleteness theorem, which he attributes Gödel as the discoverer. But this is the only thing he has learned about Gödel. Now suppose that Gödel was not the author of this theorem. A man called "Schmidt" whose body was found in Vienna under mysterious circumstances many years ago, actually did the work in question... When John uses the name 'Gödel, is he talking about:

A. the person who really discovered the incompleteness of arithmetic (descriptivist)?

or

B. the person who...claimed credit for the work (causal-historical)?

Their data suggests that East Asian students are more likely to endorse descriptivist accounts, while American students are more likely to give causal-historical responses. They consider some approaches for how referentialists might account for the data, each of which having negative consequences for proponents of arguments from reference. They could cease appeal to the method of cases in favor of some other consideration, though that would leave them without independently motivating grounds for their theory. They could endorse referential pluralism, which would lead them into the absurd situation in which they were unable to determine to what intuition group they belonged and thus whether they agreed or disagreed with one another. They could abandon substantive theories of reference outright, which would necessitate that they abandon arguments from reference (Machery, et al. 2009).

Recall that the mark of a normative project in the strong sense is the teleological prioritization of practical knowledge over theoretical knowledge. Any knowledge-that must be offered as support of or in service to knowledge-how. Machery, et al. conclude that there is convincing evidence undermining arguments from reference. They assert that referentialists must either abandon their referentialism, or else continue on knowing that there is probably substantial empirical evidence undermining their claims. One can classify their argument as a prescription: referentialism should be abandoned. This prescription is given against a backdrop of extensive engagement with the literature on arguments from reference. Presumably if there were a way to accommodate their data without undermining the necessary commitments of arguments from reference then their prescription would not stand. They are sensitive to the fallibility of their data. The authors think there probably is no such way to accommodate it other than the ways they have suggested, and so advocate the abandonment of referentialism as a “safer bet” (2009, 348). Put into the language of intellectual action and practical knowledge, for someone who is undecided with regard to how they should approach referentialism and whether it is a tenable theory of meaning, Machery, et al.’s work offers guidance. It is even easier to see how the project meets the normative expectation under the weak normative conception. The authors orient their project to examine the practical implications of their findings (that one is faced with a range of undesirable conclusions unless she rejects substantial theories of reference). They justify their recommendation that one abandon referentialism by appealing to these implications.

Machery et al.’s work and others like it are examples of experimental philosophy that do conform to the normative conception. A neo-Sellarsian philosopher might not

have qualms considering them capable of facilitating intellectual action within the philosophical landscape. It is important to note that if Machery, et al.'s project is the paradigm of experimental philosophy, then it no longer makes sense to refer to X-phi as a novel movement, as many have done. The instance of many philosophers conducting their own experiments may be sociologically novel, but it has precedent. Rose and Danks cite examples like Hermann von Helmholtz's *Treatise on Physiological Optics*, which was framed explicitly as a test of Kant's categories (2013). Such projects seem to fit Rose and Danks' broad view, provided we stipulate that *philosophical* naturalism includes a normative or theoretical component. But what should be said of the other experimental projects, those projects that neglect normative criteria and decline theoretical influence? I contend that even though such projects do not fit the normative criteria, they may yet be considered useful to the normative philosopher. In other words, they are a part of the philosophical landscape; they just do not facilitate intellectual action within it. One will still need to do some intellectual work, aided by other sources, to discern a tenable course of action.

Section 4: Philosophical Usefulness

For one thing, it seems futile to argue about whether any or all X-phi projects belong in philosophy journals (they have been and continue to be published in them), whether experimentalists are real philosophers (they have made their careers in the discipline), or whether X-phi should be taught to students of philosophy.¹⁴ But we can say something about how a normative philosopher may treat the large number of X-phi projects that do not meet the normative criteria.

Experimental philosophical projects are distinguished by their use of experimental methods. X-phi is purported to be a novel movement in philosophy because experimentalists carry out experiments or work in collaboration with those who do (Knobe 2007). So it follows that every instance of experimental philosophy includes an instance of experimentation. There is nothing inherently normative in experimentation. Experiments can serve many purposes, but they are not intrinsically normative. In order to incorporate a normative element into an experimental project, one must introduce a non-empirical component to the project.

But there is a powerful sense in which these non-normative X-phi projects seem to be at a disadvantage in discussions of their own legitimacy: it is methodologically

¹⁴ The argument has been raised by Herman Cappelen, and indirectly supported by the work of Timothy Williamson (Williamson 2007), that philosophers should reject X-phi on the basis of the unjustifiability and philosophical irrelevance of intuitions. I will not offer an argument either for or against intuitions. I do think that Cappelen's critique is inconclusive. But recall that the formulation I chose for the methodology of X-phi emphasized the examination of cognitive factors broadly construed rather than merely intuitions. This is because I take it to be fairly obvious to anyone examining the X-phi literature charitably that X-phi practitioners are not solely concerned with intuitions. X-phi projects can examine neural activity, behavior and patterns in behavior, and the effects of external manipulation on behavior and brain activity. None of these are straightforwardly concerned with intuitions, so Cappelen's characterization of X-phi is empirically misinformed (Cappelen 2012).

difficult, if not impossible for a someone to offer an argument for the legitimacy of X-phi *within* an experimental project. Experimentalists like Knobe, Neil Levy, Shaun Nichols, and others have offered arguments for the merits of X-phi, but when they do so, they are constructing arguments in a framework other than that which they are trying to defend. They can of course situate their experiments in a larger theoretical or philosophical context, but if they do so they cannot remain novel. X-phi projects cannot reflect on the role of X-phi in the larger philosophical landscape using an experiment, and the supplementation of philosophical reflection with experimental is not new. Moreover, it requires the invocation of the biasing frameworks Knobe abjures. Experimentalists under the normative conception must philosophize with the same resources of all other philosophers: a working body of theoretical knowledge, used to establish a practical theory of what philosophy is and what is relevant when trying to navigate the philosophical landscape, i.e., they use practical reason.

The experimental project that is not situated in the context of philosophical discussions, refuses to invoke existing frameworks, and fails to include a normative element will not only not meet the criteria for the normative picture, but seems to distinguish itself as a novel movement only by divesting itself of what makes a project philosophical, possibly even beyond the normative conception. Even more troubling, it risks being indistinguishable from non-philosophical scientific projects to the normative philosopher. What then is the prognosis for these experimental projects?

As I have mentioned, X-phi has the potential to be useful to a wide range of other philosophical projects. Without attempting to give an exhaustive definition of

philosophical usefulness,¹⁵ I think it is at least demonstrable that X-phi data can be brought to bear on many normative issues within philosophy. Jesse Prinz argues that empirical methods can be integral to the answering of philosophical questions. He argues that experimental philosophers may contribute to many philosophical debates by proving via experimental methods that other philosophers have approached an issue in the wrong way from the start by making false first-order predictions, e.g., assuming that there does not exist any cognitive factors that might influence one's application of the concept "intentional (Prinz 2007)." One might reject Prinz's picture of the role X-phi has to play, but it will still be the case that X-phi data can indicate patterns in concept usage, illuminate correlations between certain beliefs among those who adopt certain languages, and so forth. Insofar as philosophers are examining issues regarding the mind, language, knowledge, and existence, X-phi does have some significant findings to contribute.

One might also debate the methodological efficacy of experimental projects in their approaches to these types of issues and there is certainly room for concern about whether many X-phi experiments are unbiased, representative, or otherwise conducted properly. But this does not obfuscate the fact that these experiments do contribute data that philosophers may find useful. It may turn out that the data was wrong, but this is something that would arguably entail conducting *more experiments* to determine. Of course there is no guarantee that all X-phi projects are philosophically useful. This is similar to saying that there is no guarantee that all experiments are good ones. But something does not need to be *correct* in order to be useful.¹⁶ Newtonian physics is

¹⁵ That task might require an entire new project.

¹⁶ Joseph Pitt provided this analogy. (<http://www2.jpl.nasa.gov/basics/bsf3-2.php>)

explanatorily surpassed by the theory of special relativity, but NASA engineers still use Newton's theory to send shuttles to the moon. Moreover, even if a philosopher ultimately rejects a set of data, she might find that data useful insofar as it played a role in eliminating a bad conclusion.

Conclusion

It is tempting to conclude by saying that non-normative X-phi projects are philosophically useful but should not be understood as being properly philosophical under the normative conception. This may be a crude way of summarizing the argument I have made, but it belittles the role X-phi might have to play in the philosophical landscape. There is nothing internal to the methodology of conceptual analysis, the establishment of an ethical theory, the defense of an ontological argument, or an examination of a theory of knowledge that precludes the project's ability to include a normative component without jeopardizing the novelty of the subject. I have argued that there is such an obstacle for non-normative experimental projects. Experiments methodologically cannot be reflective on their role in philosophy without the introduction of armchair elements. The inclusion of such elements will allow the project to be differentiated from a scientific experiment under the normative conception, but will compromise its novelty. The experimentalist may not be concerned by the risk of blurring the line between philosophy and science. Or she might be willing to sacrifice novelty for that distinction. She may not be concerned to appease the normative philosopher, though I have given some reasons why she might want to do so. I will refer back to Rose and Danks to agree with their characterization of X-phi as being something other than a movement (Rose and Danks 2013). This observation is spot-on under the normative conception. X-phi should be understood as a method. So then rather than saying one is an "experimentalist," it may be more sensible to say that X-phi is something philosophers can do. Philosophers can use experimental methods to produce normatively oriented projects, or they can conduct experiments with other aims in mind. The products of non-

normative X-phi projects populate the philosophical landscape, the logical space of reasons. They are philosophical raw material. Construed in this way, there is something incredibly interesting going on in the proliferation of X-phi projects, both in and out of the normative conception.

References

- Cappelen, Herman (2012). *Philosophy Without Intuitions*. Oxford: Oxford University Press.
- Franklin, Allan (2007). "The Role of Experiments in the Natural Sciences: Examples from Physics and Biology." *General Philosophy of Science: Focal Issues*. T. A. F. Kuipers, Amsterdam: Elsevier.
- Knobe Joshua (2014). "Experimental Philosophy is Cognitive Science," in Justin Sytsma & Wesley Buckwalter (eds.) *A Companion to Experimental Philosophy*. Blackwell Publishing.
- Knobe, Joshua (2003). "Intentional Action and Side-Effects in Ordinary Language." *Analysis* 63 (3):190–194.
- Levy, Neil and Kitano, Yasuko (2011). "We're All Folk: An Interview with Neil Levy about Experimental Philosophy and Conceptual Analysis." *Annals of the Japan Association for Philosophy of Science*: Vol.19. 87-98.
- Machery, Edouard; Mallon, Ron; Nichols, Shaun & Stich, Stephen (2009). "Against Arguments from Reference." *Philosophy and Phenomenological Research* 79 (2): 332 - 356.
- Prinz, Jesse J (2007). "Empirical philosophy and Experimental Philosophy," in Joshua Knobe & Shaun Nichols (eds.), *Experimental Philosophy*. Oxford: Oxford University Press.
- Quine, W. V. (1969). Epistemology Naturalized. *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Rose, David & Danks, David (2013). "In Defense of a Broad Conception of Experimental Philosophy." *Metaphilosophy* 44 (4): 512-532.
- Ryle, Gilbert (1949). *The Concept of Mind*. New York: Hutchinson.
- Sellars, Wilfrid (1991). "Philosophy and the Scientific Image of Man." *Science, Perception and Reality*. Atascadero, CA: Ridgeview Publishing. Originally published 1963.
- Williamson, Timothy (2007). *The Philosophy of Philosophy*. Malden, MA: Blackwell Publishing.