Empirical Investigations:
Reflecting on Turing and Wittgenstein on Thinking Machines[1]

Jonathan Livengood and Justin Sytsma

In the *Blue Book*, Wittgenstein raises the question of whether it is possible for a machine to think. He writes:

> There is an objection to saying that thinking is some such thing as the activity of the hand. Thinking, one wants to say, is part of our 'private experience'. It is not material, but an event in private consciousness. This objection is expressed in the question: "Could a machine think?" I shall talk about this at a later point, and now only refer you to an analogous question: "Can a machine have toothache?" You will certainly be inclined to say: "A machine can't have toothache". All I will do now is to draw your attention to the use which you have made of the word "can" and to ask you: "Did you mean to say that all our past experience has shown that a machine never had toothache?" The impossibility of which you speak is a logical one. (1958, 16)[2]

---

[1] Penultimate draft, to appear in *Turing and Wittgenstein on Mind and Mathematics*, D. Proudfoot (ed.), OUP.

[2] It is worth noting that people (native English-speaking adults) are often willing to ascribe pains to some machines. While people are generally unwilling to ascribe pain to a simple, non-humanoid robot with hard and metallic body parts (Sytsma and Machery, 2010; Sytsma and Machery, 2012), this shifts to mild willingness if the robot's "hands" are instead made of soft and fleshy material (Sytsma, 2012). This suggests that people tend to treat pains as instantiated in body parts (see Sytsma, 2010a; Reuter et al., 2014; Kim et al., forthcoming; Sytsma and Reuter, forthcoming), and that the issue is less with the cognitive faculties of machines than with whether they are made of materials that can support pains. In contrast, people are generally willing to ascribe visual perceptual states to a simple robot (for an overview of this research see Sytsma, 2010b and 2014). Further, people seem willing to describe some psychological predicates to computer systems (Sytsma, 2010c). Of course, it might be that the use of such terms has changed dramatically since the time Wittgenstein was writing. As Turing (1950, 442) suggests, "I believe that at the end of the century the use of words and general educated opinion will have altered so much that one will be able to speak of machines thinking without expecting to be contradicted." To get a sense of popular sentiment today we asked a set of participants Wittgenstein's question—"Could a machine think?"—with participants answering by selecting either "Yes" or "No." In addition, we presented separate sets of participants variations in which "machine" was replaced by "computer" or "robot." After each question, we asked participants to explain their answers, and on a second page, we asked them either whether computers are machines (first two sets) or whether robots are machines (third set). In addition, we collected data on demographic questions. Responses were collected online through philosophicalpersonality.com and were restricted to native English-speakers, 16 years of age or older, who completed the survey and had not taken a survey through the website before. We found that 17/65 (26.2%) answered affirmatively when asked whether a machine could think. This percentage rose when participants were asked whether a computer could think, with 27/75 (36.0%) answering affirmatively, although the difference was not statistically significant ($\chi2=1.1429$, $p=0.285$ two-tailed). And rose again when participants were asked whether a robot could think, with nearly half answering affirmatively (27/56, 48.2%). Although the difference between the computer question and the robot question was not statistically significant ($\chi2=1.5022$, $p=0.2203$, two-tailed), the difference between the machine question and the robot question was ($\chi2=5.4092$, $p=0.02003$, two-tailed). Despite this, nearly everyone answered that computers are machines (138/140, 98.6%) and everyone answered that robots are machines (56/56). These results will be discussed at greater length elsewhere.

When Wittgenstein returns to the issue of a thinking machine, he elaborates on the suggestion that the question concerns logical possibility as opposed to empirical possibility:

> "Is it possible for a machine to think?" …. And the trouble which is expressed in this question is not really that we don't yet know a machine which could do the job. The question is not analogous to that which someone might have asked a hundred years ago: "Can a machine liquefy a gas?" The trouble is rather that the sentence, "A machine thinks (perceives, wishes)": seems somehow nonsensical. It is as though we had asked "Has the number 3 a colour?" (47-48)[3]

One thread to follow from this passage is the claimed disanalogy between the question of whether a machine can think and the question of whether a machine can liquefy a gas. A straightforward interpretation is that Wittgenstein holds that whether a machine can think is a *conceptual* matter, but whether a machine can liquefy a gas is a *scientific* or *empirical* matter. It might then be thought that the answer to the first question can be found without employing the methods of the sciences.

A more nuanced reading is that Wittgenstein holds that the first question raises a "philosophical problem" while the second question does not. Wittgenstein seems to feel some urge to say that talk of a thinking machine sounds absurd, despite harboring no such urge with regard to talk of a machine producing the outward signs of thinking. And it is this disconnect that appears to grab his interest. Such a reading is supported by the context in which Wittgenstein first raises the question of a thinking machine in the *Blue Book*. This occurs in response to the objection that it sounds absurd to say that thinking is an activity of the hand, but rather takes place in the mind. The immediate context is that the statement "thinking is a mental activity"—and the analogous question "Where does thinking take place?"—is liable to mislead us through "misunderstanding the grammar of our expressions" (15-16). The route of misunderstanding appears to be in thinking that there must be some one locality for thought, such that "the *real*

---

[3] A similar passage is found in *Philosophical Grammar* (1974, 105).

place of thought is in our head" or in private consciousness (16). Put another way, Wittgenstein suspects that the question of whether a machine can think is liable to bewitch us in a way that the question of whether a machine can liquefy a gas is not. And for Wittgenstein, "philosophy is a battle against the bewitchment of our intelligence by means of language" (1958, §109).

On both readings, we find that Wittgenstein's focus is on "conceptual" rather than "empirical" questions.[4] In fact, just a couple of pages after the quote from the *Blue Book* that we opened with, Wittgenstein gives a famous passage casting aspersions on the use of the method of science in philosophy:

> Philosophers constantly see the method of science before their eyes, and are irresistibly tempted to ask and answer questions in the way science does. This tendency is the real source of metaphysics, and leads the philosopher into darkness. I want to say here that it can never be our job to reduce anything to anything, or to explain anything. Philosophy really *is* 'purely descriptive'. (18)

Passages such as this have led many interpreters to view Wittgenstein as taking his approach to philosophy to be utterly distinct from the scientific approach. For instance, in his introductory remarks on Wittgenstein's life and work, Sluga (1996, 25) writes that "Wittgenstein insists generally on a sharp distinction between philosophy and science, in sharp contrast to those movements in the twentieth century that have sought to reconstruct philosophy in a scientific manner." And after reviewing passages like that given above, Sluga notes that "to philosophers steeped in the values of science such remarks must naturally sound offensive" (25).

Not too surprisingly, the received wisdom has it that "for Wittgenstein philosophy is nothing like science" (Garver, 1996, 151). We, on the other hand—as advocates of experimental

[4] As Wittgenstein later notes in *Culture and Value* (1980, 79e): "I may find scientific questions interesting, but they never really grip me. Only *conceptual* and *aesthetic* questions do that. At bottom I am indifferent to the solution of scientific problems; but not the other sort."

philosophy[5]—think that philosophy is often advanced by the application of "scientific" tools to "philosophical" problems. Yet, the experimental philosophy we advocate is derived, in part, from our reading of Wittgenstein. And we believe that even if we restrict the enterprise to conceptual questions (or philosophical problems, in Wittgenstein's sense), there is much to be gained by empirical investigation.[6] Thus, when Wittgenstein writes in §109 of the *Philosophical Investigations* that "philosophical problems… are, of course, not empirical problems; they are solved, rather, by looking into the workings of our language, and that in such a way as to make us recognize those workings," we want to respectfully disagree: we believe that such problems are *also* empirical problems, for we believe that many empirical methods are well suited for helping us look into the workings of our language. To draw out how one might arrive at an experimental philosophy by way of Wittgenstein, we will focus on the question of whether machines can think. And we'll begin with probably the most famous paper on thinking machines: Alan Turing's (1950) "Computing Machinery and Intelligence."

1. Thinking Machines

Turing opens his seminal paper on "Computing Machinery and Intelligence" in a way that we, as experimental philosophers, find tantalizing and provocative. He writes:

---

[5] For an introduction to experimental philosophy see the texts by Alexander (2012) and Sytsma and Livengood (2015). For collections of articles in experimental philosophy, see Knobe and Nichols (2008, 2013), as well as the volumes in the *Advances in Experimental Philosophy* series from Bloomsbury Publishing and series editor James Beebe. For discussions of recent disputes in experimental philosophy, see Machery and O'Neil (2014). And for an extensive survey of the state of the art of experimental philosophy, see *A Companion to Experimental Philosophy* edited by Sytsma and Buckwalter (2016).

[6] We are not inclined to restrict the discipline in this way and believe that there are many problems that philosophers can fruitfully engage with that not clearly just conceptual questions. In line with this we embrace a broad conception of experimental philosophy that extends beyond the investigation of concepts or language (Sytsma and Livengood, 2015). While many experimental philosophers have described the field in narrow terms, a growing number have either argued for or explicitly endorsed a broad conception (Sytsma and Machery, 2013; Rose and Danks, 2013; O'Neill and Machery, 2014; Sytsma and Livengood, 2015; Weinberg, 2015; Buckwalter and Sytsma, 2016; Stich and Tobia, 2016; Schupbach, forthcoming). For a recent discussion of the narrow versus broad conceptions see Sytsma (forthcoming).

> I propose to consider the question, ‘Can machines think?’ This should begin with definitions of the meaning of the terms ‘machine’ and ‘think’. The definitions might be framed so as to reflect as far as possible the normal use of the words, but this attitude is dangerous. If the meaning of the words ‘machine’ and ‘think’ are to be found by examining how they are commonly used it is difficult to escape the conclusion that the meaning and the answer to the question, ‘Can machines think?’ is to be sought in a statistical survey such as a Gallup poll. But this is absurd. Instead of attempting such a definition I shall replace the question by another, which is closely related to it and is expressed in relatively unambiguous words. (1950, 433)

What is most tantalizing and provocative about this passage, for us, is the charge that conducting a survey would be “absurd.” The question is why exactly this is absurd (if indeed it is)? One possibility, and one that we are swayed by, is that such a survey *on its own* is unlikely to settle the question, for it is likely that we will also need to know what the theoretical capacities of machines are. The bulk of the absurdity, on this account, would then be due to the specific question being asked.

The surrounding discussion in this passage suggests that this is not quite Turing’s worry, however, for he begins by asserting that while we need to define the terms “machine” and “think” it would be dangerous to simply focus on the normal use of the terms. And the absurdity of the Gallup poll is seemingly called on as support for this claim. Turing then proposes that rather than define these terms in line with ordinary usage, he will replace the question with one that is “closely related to it and is expressed in relatively unambiguous words.” But how closely related does the new question need to be for it to serve as an adequate replacement? And what is the criterion for adequacy in this type of investigation? The worry here is that replacing the question might simply be to *change the question*, abandoning it and proceeding with a wholly different investigation. Plausibly, the adequacy of the replacement will depend on just how “closely related” it is, but it is not obvious how we are to determine how close a proposed replacement is if we do not first get clear on the meanings of “machine” and “think.” But how do we make our ideas of machines and thinking clear?

In his *Philosophical Investigations*, Wittgenstein famously said that “For a *large* class of cases—though not for *all*—in which we employ the word ‘meaning’ it can be defined thus: the meaning of a word is its use in the language” (§43). But the idea that meaning is use was not new to the *Investigations*. Wittgenstein was already developing and defending the idea in *The Blue and Brown Books*, dictated between 1933 and 1935, and in lectures he gave on the foundations of mathematics in 1939. In the first of his lectures on the foundations of mathematics, Wittgenstein (1976, 19) said that “to understand a phrase, we might say, is to understand its use.” In later lectures, he consistently directed the members of his seminar to consider the use of mathematical symbols, the practical goals of calculations, and the way mathematical techniques are taught and learned. But if Turing was right in thinking that the thesis that meaning is use has the consequence that one should seek the meaning and the answer to the question “Can machines think?” with a Gallup poll, and if Turing was right that such a proceeding would be absurd, then Wittgenstein’s thesis leads to absurdity and (presumably) ought to be rejected.

Turing could not have been ignorant of Wittgenstein’s view that meaning is use. We know that Turing attended many of Wittgenstein’s 1939 lectures on the foundations of mathematics, and at least in some instances, students recorded what Turing said during the sessions. In his first lecture, Wittgenstein mentioned Turing five times, in the playful way one sometimes does when a friend is in the audience. In his fourth lecture, Wittgenstein said that in calculating a weight, there are two different ways one might use an arithmetical proposition such as $25 \times 25 = 625$. The proposition is experiential if it is used to predict the result of some physical measurement. For example, if we were to lay out 25 rows with 25 objects in each row, we might predict that if each object weighs one gram, then the entire collection weighs 625 grams. Used in this way, Wittgenstein suggests that the proposition would be false if the actual measure of the entire collection were different from 625 grams. Perhaps if each individual mass were imperceptibly heavier than one gram, the entire collection would weigh more than 625 grams. But in another (more typical) sense, the proposition that 25 × 25 = 625 is independent of experience "because nothing which happens will ever make us call it false or give it up" (41).

Wittgenstein continued with another illustration in terms of length. Suppose one were to hold up a piece of chalk and say, "This piece of chalk is the unit of length." One might mean to be making an empirical claim to the effect that all (actual) objects have lengths equal to a natural number of chalk-lengths. Or one might mean to be establishing a rule for measuring lengths. Turing asked whether the second option is a case of definition, and Wittgenstein responded by asking what is on the right-hand side and what is on the left-hand side of the definition. He continued (42):

> One might write "This piece of chalk = 1 W df." But although one writes down the words and puts the equality sign and so on, one does not write down how the definition is to be used. The important thing would be teaching a man the technique of measuring with this unit. It is no good simply writing "This piece of chalk = 1 W df." One has to say, "Take this piece of chalk and put it alongside the object in question, and then. . .", teaching him the technique. We often put rules in the form of definitions. But the important question is always how these expressions are used.

Turing could not have missed the emphasis that Wittgenstein placed on paying attention to how language is used. It is very natural, then, to read Turing's remarks at the beginning of "Computing Machinery and Intelligence" as a slap at Wittgenstein and ordinary-language philosophy. Wittgenstein's dangerous philosophy has absurd consequences.

Turing's claim that it would be absurd to use a statistical survey to look for the meaning and the answer to the question "Can machines think?" initially appears rather obvious and unobjectionable. But then—as Turing recognized—to many people, the claim that machines cannot do many of the things essential to thinking, such as being kind, telling right from wrong, learning from experience, or doing something really *new*, also appears rather obvious and unobjectionable. Turing canvassed these and other objections to the claim that machines can think. The argument appears to be as follows:

> [1] No machine can do X.
> [2] Everything that can think can do X.
> ---
> [3] No machine can think.

Turing conjectured that people endorse premise [1] on the basis of an inductive argument. He writes (447):

> A man has seen thousands of machines in his lifetime. From what he sees of them he draws a number of general conclusions. They are ugly, each is designed for a very limited purpose, when required for a minutely different purpose they are useless, the variety of behaviour of any one of them is very small, etc., etc. Naturally he concludes that these are necessary properties of machines in general.

In Turing's phrasing, the machines common to ordinary experience have "very small storage capacity." Consequently, people incorrectly infer that the limitations of small-storage machines are limitations of machines in general. An alternative hypothesis is that people do not have a good sense for how the range of possible behaviors of a machine scales with the machine's complexity. Hence, people do not see how something as flexible as human behavior could be produced by a machine.

Similarly, using a statistical survey to look for the meaning and the answer to the question "Can machines think?" seems absurd (in part, at least) because what one imagines when one thinks of a statistical survey is not a suitably complex empirical instrument. The ordinary use of the words "machine" and "think" probably cannot be discovered by way of a simple questionnaire that directly asks "Can machines think?"[7] But if we take a broad view of what counts as a statistical survey, including true experiments in which features of the instrument are manipulated across treatment groups, then the project is not absurd.[8] Further, one might investigate the ordinary uses of "machine" and "think" by way of reaction time studies, semantic integration studies, memory and false-recall studies, belief-updating studies, neuroimaging studies, simulation and computational studies, corpus analyses, and so on.

We believe that one could investigate the ordinary meanings of the terms "machine" and "think" using the tools mentioned above without committing to the thesis that meaning is use. But let us provisionally accept this Wittgensteinian dictum. On that supposition, we hold that the meanings of the terms "machine" and "think" can be fruitfully investigated by using something like a "statistical survey"—that is, by employing empirical methods commonly used in the social sciences. In fact, we believe that if one takes the ordinary meanings of "machine" and "think" to be relevant to understanding the question of whether machines can think, then *some* sort of empirical investigation will be required. But simply asking people a question like, "Can machines think?" directly need not be very illuminating (although the results might be surprising, as might the variation in responses across groups).[9] What people say in response to a direct

[7] It is plausible that there is not one single ordinary use of these terms, even restricting ourselves to English and their use in contexts like those of asking whether a machine can think. In fact, the results of the simple study we described in Footnote 2 suggest that this is likely the case, with dissenting minorities holding that a machine could think (26.2%), that a computer could think (36.0%), and that a robot could think (48.2%). It is a tricky question what we should take from such disagreement and what implications it has, and one that we will not tackle here. Given such disagreement, however, it is probably more accurate to speak of either ordinary uses (plural) or of the dominant use of such terms, again recognizing that this might well be context dependent.

[8] See Part II of Sytsma and Livengood (2015) for a discussion of questionnaire methods, including their use in both descriptive surveys and true experiments.

[9] To illustrate, we find the disagreement with regard to each of the primary questions asked in the study described in Footnote 2 to be quite interesting. We find a notable minority answering that a machine could think, that a computer could think, and that a robot could think. Also interesting, we find that the size of this minority grows across these questions, despite an overwhelming majority answering that computers are machines and that robots are machines. This might suggest that some answers to these questions are being driven by rather superficial features of the questions, perhaps that it is more common to switch to terms like "computer" or "robot" today in the context of intelligence or cognition, than to talk about machines. This finding might offer some very preliminary support for

question does not tell us *why* they answer in the way they do. Suppose someone answers that machines cannot think. Is she correctly deducing her answer from her concepts or is she—as Turing suggested—making a mistaken inductive inference? A better approach would be to ask several questions and to conduct experiments by manipulating various parts of the stimuli presented to participants in order to identify what people mean by "machine" and "think." Then one could ask whether people *ought* to say that machines can think given the way they use the individual terms.

Notice that even if we could completely determine the ordinary meanings of the terms "machine" and "think" by empirical investigation, we might or might not thereby settle the question of whether machines can think. An empirical investigation into the meanings of the terms might settle the question if—in line with the theological objection that Turing considered (443-444)—ordinary people took *being a machine* to rule out having a soul and ordinary people took *being able to think* to necessitate having a soul. On such conceptions, it is a deductive consequence that no machine can think.[10] Alternatively, suppose people thought that anything is

---

the type of worry we voiced for Turing above, and could be further explored by using other methods such as corpus analysis. Finally, we noted that Turing believed that opinions about thinking machines would shift dramatically by the end of the previous century. In this context, it might be thought to be rather surprising that only about one in four of our participants answered that a machine could think. Further, based on this suggestion we might expect to see an age effect in responses, with older participants being more likely to deny that a machine could think than younger participants. This is not what we found, however, with there being a very slight positive correlation between affirmative responses and age (r=0.03881). Alternatively, focusing on Turing's phrase "general educated opinion," we might expect to see a positive correlation between education level and participant responses. This time we did find a notable correlation (r=0.21004), with more highly educated participants being more likely to give affirmative responses. Of those participants who had completed some college or more, 10/29 (34.5%) gave an affirmative response compared to just 7/36 (19.4%) of the remaining participants.

[10] Recall that in the study discussed in Footnote 2, in addition to asking participants whether a machine could think, we asked them to explain their responses. While such qualitative data can be difficult to interpret, it is often informative. While a careful analysis of such data requires independent coding, it is worth noting that several explanations indicated that some participants were operating with a conception in the vicinity of the one that Turing suggests. For instance, a number of participants noted that machines lack brains and that brains are necessary for thinking (i.e., "a machine doesn't have a brain, so therefore it cannot think"), even as a few others wrote that a machine could think because it has a brain (i.e., "a computer has a brain, so there should be some kind of connection to thinking I guess"). Others noted similar connections to the need for feelings, emotions, consciousness, mind, free will, or simply being alive or human. Given such explanations, we might expect that there will be a correlation between participants' beliefs and how religious they are (especially recognizing that the sample is predominantly

a machine if its behavior is well-described as "calculating" and that thinking is well-described as calculating. On such conceptions, *anything* that thinks is a machine. So, if anything can think, a machine can think.[11] If the point of asking whether machines can think is to decide what one should say given the ordinary conceptions of "machine" and "think," then the question could be answered by an empirical investigation into the meanings of the terms, provided the answer to the question is a logical consequence of the ordinary conceptions.

But, of course, it could turn out that the answer is not a logical consequence of the ordinary conceptions. Suppose that according to the common usage of "machine," any composite object that a human can construct by combining and arranging tokens of the six types of simple machines (levers, pulleys, wheels, inclined planes, wedges, and screws) is a machine. And suppose that according to the common usage of "thinking," something is capable of thinking if it can hold and express desires, observe and represent its environment, draw conclusions by inference from its representations, generate lists of possible actions for achieving its desires given arbitrary environments, select actions from the lists it generates, and modify its desires, representational rules, inferential rules, or decision rules if it fails to achieve its desires. Given those conceptions, whether a machine can think is an empirical question of a different kind. To ask whether a machine can think is to ask whether it is possible for a human to arrange tokens of

---

drawn from Americans). And, in fact, we found a sizable correlation between religiosity (defined in terms of participating with an organized religion) and how likely people were to answer that a machine could think: the more religious a person was the less likely they were to give an affirmative response (r=-0.3398). Of those participants answering "Not at All" or "Almost Not," 15/36 (41.7%) gave an affirmative answer, compared to just 2/29 (6.9%) of the remaining participants.

[11] A few respondents to the survey described in Footnote 2 offered explanations in this vicinity (i.e., "In a way, all machines think. By processing a series of code, they reach conclusions. The human mind operates in a similar manner, differing primarily in the level of complexity…").

various simple machines into a single composite object that can perform the actions that characterize thinking things.[12]

A third possibility is that the use of a statistical survey to answer the question of whether machines can think is absurd because it has the absurd consequence that philosophy is as easy as carrying out a survey sample.[13] Like Swift's professor in the Grand Academy of Lagado who aimed to improve speculative knowledge "by practical and mechanical operations," one might think that with empirical methods, "the most ignorant person, at a reasonable charge, and with a little bodily labor, might write books in philosophy … without the least assistance from genius or study" (51). We agree that philosophy is not easy. But, then, neither is survey sampling. Moreover, it's not clear how philosophy as a *whole* would be made easy if we were able to make *some* advances in philosophy by way of a Gallup poll.

Alternatively, one might argue that the use of empirical methods is absurd because it makes the absurd assumption that the person on the street is never wrong. But there are two distinct ways that the person on the street might go wrong in answering the question, "Can a machine think?" First, she might answer the question in a way that is wrong by the lights of her own typical usage. For example, she might generally use the terms "machine" and "think" in a way that allows that machines could think and yet confidently assert that they cannot or vice-versa. After all, it is surely not just philosophers who can be bewitched by language. Alternatively, she might use the terms so as to make the question an empirical one, and yet confidently assert that it is impossible for machines to think. Or her expressed views about the question might come apart from her own usage in some other way. Second, she might be

[12] Again, a number of respondents gave explanations that suggested that they saw this as an open empirical question (i.e., "I have seen things happen in my lifetime that I never thought possible. Therefore, this may also be possible.").
[13] We think Turing deploys essentially this argument against Wittgenstein's view of mathematics as invention, rather than discovery, an issue that we'll return to in Section 2.

incorrectly using the terms “machine” and “think.” One might be tempted to deny that it is possible to go wrong in this way. Surely there are no meanings, residing in Plato’s heaven, against which we check the use of our terms! However, one need not have an inflated ontology in order to make sense of the thought that a person can mean the wrong things by her terms. If a person intends to be guided by some external source, such as a linguistic community or a relevant expert, then she might use her terms incorrectly insofar as her use of those terms deviates from the use prescribed by the relevant standard. For example, when Constable Dogberry in *Much Ado About Nothing* says, “Our watch, sir, have indeed comprehended two auspicious persons,” he is using “comprehended” and “auspicious” incorrectly. More interestingly, if a person is (implicitly or explicitly) using a term for some purpose, then she could use it incorrectly by using it in a way that fails to serve that purpose.

Hence, regardless of how the terms “thinking” and “machine” are related to each other in the relevant linguistic community, empirical work has a role to play. If empirical work on the meanings of the terms shows that “thinking” and “machine” are related in the right way, then one does not need to build a machine or prove a theorem in order to say whether a machine can think. On the other hand, if empirical work on the meanings of the terms shows that neither term entails the other (or its negation), then further empirical work is required in order to say whether a machine can think. One would need to build a machine that satisfies the requirements of thinking or prove an impossibility theorem that rests on plausible physical axioms. Moreover, as we will see later, if one is inclined (as Turing was) to replace the question of whether a machine can think, then empirical work on the terms “thinking” and “machine” might be useful for establishing that the replacement question is both closely-enough related to the original question and also interesting enough to warrant careful consideration.

2. Bolshevism in Mathematics

Turing complained that the question of whether machines can think is “too meaningless to deserve discussion” (442), which resonates with Wittgenstein’s remark that to ask whether a machine thinks seems nonsensical in the same way that it seems nonsensical to ask whether a number has a color. But there is an important difference. Turing took the question to be meaningless in a way that would admit *replacing* it with a closely-related but clearer alternative. Wittgenstein, by contrast, thought the question was meaningless in a less tractable way. For plausibly, there is no closely-related but clearer alternative to the question of whether a number has a color. The question is simply irredeemable. We take the difference here to be one instance of a general disagreement between Turing and Wittgenstein about philosophical method and philosophical progress. In this section, we briefly characterize the disagreement. We then discuss another instance of the disagreement between the two having to do with whether advances in mathematics are a matter of invention or discovery. Finally, we come back to the question of whether machines can think.

Turing suggested replacing the question of whether a machine can think with an alternative based on *the imitation game*. In the imitation game, a human and a machine sit out of view of an interrogator, who communicates with them by sending and receiving text messages through an independent computer system. The interrogator tries to identify which of the two is the machine, and the reliability with which the interrogator does so is a measure of how successfully the machine plays the game. Turing replaces the question of whether machines can think with the question of “whether there are imaginable computers which would do well” at the imitation game. Turing’s replacement question *does* seem to be analogous to the question, “Can a machine liquefy a gas?” If we exclude humans from consideration—perhaps being in doubt as to whether they are really *machines*—then we do not know whether a machine can be constructed that would do well at the imitation game.[14]

But for Wittgenstein, one does not make philosophical progress—or gain philosophical insight—by replacing a question that raises a philosophical problem with a more tractable question. If one replaces a puzzling question *in philosophy*, the philosophical problem remains. One has simply chosen not to talk about it. The work of philosophy consists in uncovering the sources of our philosophical problems, and in seeing what was puzzling about the original question in the first place. Wittgenstein has clear ideas about the general sources of our philosophical problems, as he remarks in §119 of the *Philosophical Investigations*: "The results of philosophy are the uncovering of one or another piece of plain nonsense and of bumps that the understanding has got by running its head up against the limits of language."[15]

For Wittgenstein, philosophy does not interfere with the use of language, but only describes it—"it leaves everything as it is" (§124). And the same holds for mathematics.[16] As Wittgenstein continues in §125:

> It is the business of philosophy, not to resolve a contradiction by means of a mathematical or logico-mathematical discovery, but to make it possible for us to get a clear view of the state of mathematics that troubles us: the state of affairs *before* the contradiction is resolved. (And this does not mean that one is sidestepping a difficulty.)

We think that in the cases that Wittgenstein had in mind, logico-mathematical discovery serves to resolve a contradiction in one sense, but not in another. As Turing suggests, there is a sense in which we resolve a contradiction by replacing some notions with others. The contradiction is

[14] We (the authors) have no live doubts that a machine can be constructed to pass the Turing test: a machine that passes the Turing test *will* be built; it is only a matter of time.
[15] Such results can then serve as reminders, helping us to avoid mistakes in the future. As Wittgenstein writes in §127, "The work of the philosopher consists in assembling reminders for a particular purpose." The thought here is perhaps similar to Peirce's observation (1878, 301) that "Metaphysics is a subject much more curious than useful, the knowledge of which, like that of a sunken reef, serves chiefly to enable us to keep clear of it."
[16] "It also leaves mathematics as it is, and no mathematical discovery can advance it." (§124).

resolved to the extent that after our revisions to the language, we know how to go on. But in another sense, the contradiction is never resolved. Mathematicians, for example, in deciding to go on in the way they do, change the meanings of their terms. But the contradiction remains with respect to the *original* meaning of the terms (the way the terms are used in the state of affairs where we have not yet decided how to go on).

Perhaps unsurprisingly, then, the same disagreement that we have been sketching also appeared in exchanges Turing and Wittgenstein had on the foundations of mathematics. In the first of his 1939 lectures on the foundations of mathematics, Wittgenstein said:

> One talks of mathematical discoveries. I shall try again and again to show that what is called a mathematical discovery had much better be called a mathematical invention. (22)

In subsequent lectures, Wittgenstein described some simple cases to illustrate and support the claim that in mathematics, we make decisions about how to extend the use of our terms into new domains—a process that is best thought of as invention. In developing the idea that mathematics is about deciding how to extend the use of our terms into new domains, Wittgenstein asked what it means to say that we *can* construct a pentagon but that we *cannot* construct a heptagon.[17] In his sixth lecture, Wittgenstein suggested that there are two senses of the terms "same," "similar," and "analogous." He gave the following example. Suppose we write down a sequence of numbers—such as 2, 4, 6, 8—and then we say to a student, "Do something analogous starting with 102." That is one sense of "analogous." Alternatively, we might say to the student, "The analogous thing to write down next is 104." According to Wittgenstein, in the first case, we are asking the student to apply a certain training about how to use the term "analogous," and in the

[17] By "construct," we mean "construct with an unmarked straight-edge and compass in Euclidean geometry." That a regular pentagon can be constructed using only straight-edge and compass was known in antiquity. In 1837, Pierre Wantzel proved that a regular heptagon cannot be constructed. See Suzuki (2008) for discussion of some geometric impossibility results.

second case, we are actually providing such training. The idea is that what the mathematician does is to invent and to teach a new interpretation of “analogous” or of “going on in the same way.” Put another way, mathematicians propose rules for saying that two things are analogous or for extending an existing practice into a new domain of application.[18]

In order to make the contrast between the two senses of “analogous” clearer, Wittgenstein returned to the case of polygon constructions (66). He imagined that we have two (qualitatively identical) drawings of a pentagon but use those drawings in different ways. We take the first drawing and use it to draw a new picture of a pentagon. By contrast, we take the second drawing and use it to “draw the construction of a heptacaidecagon.”[19] Regarding the second case, Wittgenstein remarked:

> Turing might say this is drawing something in a different projection. Is this the case of the man who invented the construction of the heptacaidecagon? Does he follow a rule for projecting it (like drawing on a different scale)? Wasn't he *introducing* a new mode of projection? He *invented* a new mode of projection, which there is reason to call so-and-so. He discovered a new kind of analogy. He had learnt one mode of projection in the one case, but not in the other. He was given a picture. And the point was to invent a mode of projection. (66)

Turing—unconvinced—replied, “It certainly isn't a question of inventing what the word ‘analogous’ means; for we all know what ‘analogous’ means.” Wittgenstein responded by saying that inventing the construction of the heptacaidecagon was indeed to give a new meaning to the word “analogous” but that inventing the construction was not *merely* about how the word “analogous” was to be used:

> Yes, certainly, it’s not a question merely of inventing what [“analogous”] is to mean. For if that were the problem, we could settle it much easier by making “analogous” mean “cushion”. The point is indeed to give a new meaning to the word “analogous”. But it is not merely that; for one is responsible to certain things. The new meaning must be such that we who have had a certain training will find it useful in certain ways. It is like the

[18] In this connection, one might consider the process of analytic continuation in mathematics.
[19] A heptacaidecagon (sometimes “heptakaidecagon” or simply “heptadecagon”) is a 17-sided polygon. In 1796, Gauss proved that the heptacaidecagon is constructible with compass and straightedge. See Weisstein (2017).

> case of definitions. Is a definition purely verbal or not? Definitions do not merely give new meanings to words. We do not accept some definitions; some are uninteresting, some will be entirely muddling, others very useful, etc. And the same with analogy. (66)

Wittgenstein continued by saying that the person who invented the construction of the heptacaidecagon need not have actually drawn anything. The inventor might have found the construction already written out and simply been the first to recognize its significance. In that case, Wittgenstein asked, what has the inventor found, since the entire construction was there already? We take Wittgenstein's question to be rhetorical. We know what Wittgenstein *wants* us to say—that the inventor hasn't *found* anything.[20] Rather, the inventor changes the meaning of "construction" to include what was found. But Turing did not quite give the answer Wittgenstein wanted. Instead, Turing answered, "He sees the analogy with the construction of the pentagon."

Unsatisfied, Wittgenstein offered a different example, in which we imagine that there is a white lion in a room full of people:

> If the lion had always been in the room it couldn't have been found. Suppose everyone had seen the white lion but hadn't realized it was a white lion. [Some fellow] suddenly realizes that this is the picture of that. But what does it come to, to say that he suddenly realizes this? He gives "white lion" a new meaning. (67)

Turing replied, "I understand but I don't agree that it is simply a question of giving new meanings to words." At this point, Wittgenstein declared victory and gave a very provocative explanation of the apparent disagreement:

> Turing doesn't object to anything I say. He agrees with every word. He objects to the idea he thinks underlies it. He thinks we're undermining mathematics, introducing Bolshevism into mathematics. But not at all. We are not despising the mathematicians; we are only drawing a most important distinction—between discovering something and inventing something. (67)

[20] In his fifth lecture, Wittgenstein mentioned an exchange with Lewy in which Lewy replied to a question by saying, "Well, I know what you want me to say." After mentioning Lewy's reply, Wittgenstein remarked (55), "That was a severe criticism of me. For I have no right to want you to say anything except just one thing: 'Let's see'."

In light of the above discussion of Wittgenstein's perspective on the goals of philosophy, this passage can be read as pointing out a disconnect between Wittgenstein's concerns and Turing's. While Turing is concerned with the implications of Wittgenstein's views for mathematics, Wittgenstein is concerned with a philosophical problem and holds that philosophy "leaves mathematics as it is." Regardless of Wittgenstein's intentions, however, it is fair to worry that his views might nonetheless have the implication of introducing Bolshevism into mathematics.[21]

What would it mean to introduce Bolshevism into mathematics (or any other discipline, for that matter)? And why did Turing worry about it (assuming Wittgenstein was correctly assessing the situation), as opposed to welcoming the revolution? We think that there are three interrelated themes one might draw out of the idea of introducing Bolshevism into mathematics. First, there is an emphasis on the community, on what is held in common, and on what is public. These are related to the idea of meaning as use in a language. If you want to know what is meant by a term or a sentence, then you need to look at how the term or sentence is used *publicly* in a *community of speakers*. Second, there is an elevation of the commonplace—the ordinary—and a correlative depression of the extraordinary, including a dismissal of experts and expertise. Third, there is a revolutionary undertone. Introducing Bolshevism means overthrowing the existing regime. And specifically, it is overthrowing *from below*. On our reading, Turing mostly objected to the second and third points. But he took the first point to provide the motive force.

We take Turing to have been raising the same basic worry in reaction to Wittgenstein's view of mathematics and in reaction to an application of Wittgenstein's philosophy to the question of whether a machine can think. To seek the meaning and the answer to the question, "Can a machine think?" by way of a Gallup poll would be to introduce Bolshevism into

[21] The same basic points of disagreement recur when Turing and Wittgenstein talk about "mathematical experiments," especially in Wittgenstein's tenth lecture. But we do not have space in this essay to elaborate.

philosophy. And as we suggested that Turing might find the idea of using a Gallup Poll to answer the question of whether machines can think to be absurd because it would make philosophy too easy, the worry about introducing Bolshevism into mathematics might be expressed by saying that it would make mathematics too easy. It would replace expert training and careful thought with populist assent. But true progress in mathematics is not so easy to accomplish as that. The truths of mathematics are not a matter to be settled at the polls.

We believe that there is a close analogy between these worries about Bolshevism and common worries about experimental philosophy over the past decade. A frequent refrain (at least in conversation) is that experimental philosophy is too easy, replacing expert training and careful thought with popular opinion. But, as in mathematics, true progress in philosophy is not so easy to accomplish as that. Not surprisingly, there is an analogue in discussions of experimental philosophy for each of the three interrelated themes noted above for Bolshevism in mathematics. Experimental philosophy has often been described in a narrow way that focuses on the empirical study of philosophical intuitions, especially the intuitions of lay-people. And there is a clear connection between this focus and investigating the public use of philosophically interesting terms among a general community of speakers. On its own, such investigations should not seem overly problematic, although some philosophers see them as irrelevant to addressing philosophical problems.[22] But many prominent experimental philosophers investigating lay intuitions about philosophical concepts have not advocated for a merely descriptive project, as Wittgenstein does, but have conducted their investigations with an eye toward replacing appeals to philosophers' own intuitions with the results of more general and systematic surveys of lay

[22] See Williamson (2016) for a recent discussion of philosophical criticisms of experimental philosophy, including the relevance of investigating intuitions. See Chapter 4 of Sytsma and Livengood (2015) for a discussion of some typical criticisms of experimental philosophy and our responses.

intuitions. And some critics have seen this as an attack on (one area of) philosophical expertise.[23] Finally, experimental philosophy has often been seen as an attempt to overthrow “traditional” philosophy from below, and the common image of the burning armchair from the early days of the new experimental philosophy supports fears about its revolutionary intentions. While we support a broad conception of experimental philosophy on which it has no special connection to the study of philosophical intuitions, and on which the experimental turn need not be motivated by or even support radicalism, such concerns are slow to die.

Coming back to the question of whether machines can think, we read Turing as trying to make progress by replacing a vague question with a more precise one. But from Wittgenstein’s point of view, Turing was changing the question, and in so doing, he was sidestepping the philosophical problem that Wittgenstein was interested in. It is not that Turing was wrong by Wittgenstein’s lights. Rather, the problem is that offering a new way to go on casts no light on the bump in the road that was blocking our path. Wittgenstein’s concern was with uncovering why we felt the urge to treat the question “Can a machine think?” as being disanalogous with the question “Can a machine liquefy gas?” despite their surface similarity. Wittgenstein was aiming to lay bare the sources of our metaphysical assumptions about thought that led to this sense of disanalogy. Turing, by contrast, was offering a replacement that would eliminate the disanalogy. We might say that Turing attempts to move beyond the conceptual worry that Wittgenstein is interested in by changing the question “Can a machine think?” to one as straightforwardly empirical as “Can a machine liquefy a gas?”

Hence, we see the disagreement between Turing and Wittgenstein, both with regard to mathematics and with regard to philosophy, as reflecting a difference in their fundamental aims.

[23] See Alexander (2017) for a recent survey of debates concerning experimental philosophy and philosophical expertise.

As we noted above, however, one might have a concern with a philosophical view not because of what it was intended to address, but because of the implications it has. Thus, Turing had a legitimate reason to worry about the potential implications of Wittgenstein's views for mathematics, even if that wasn't the purpose of Wittgenstein's investigation. Similarly, we find that there is pragmatic reason to worry about the move that Turing made on the machine question beyond simply being concerned with the philosophical problem being elided. One pragmatic worry about replacing the question "Can a machine think?" with a more straightforwardly empirical one is that insofar as it leaves the original philosophical problem untouched—and insofar as we continue to feel that a thinking machine is in some way nonsensical—the replacement question is unlikely to do the desired work. If one is inclined to think that the idea of a thinking machine is nonsensical, then one will be inclined to think either that the idea of a machine passing the Turing test is absurd or that passing the Turing test is not sufficient for thinking. The sympathy that many people have had with Searle's Chinese Room thought experiment over the past four decades illustrates just this worry. In a sense, this is a return to the suggestion from Section 1 for why Turing might find appeal to a Gallup Poll to be absurd that we noted above: Turing's strategy would make philosophy too easy. His strategy avoids the difficult battle of uncovering how our thinking about the machine question has been bewitched, and in so doing is unlikely to break the spell.

3. Concluding Remarks: Two Experimental Turns

We began this chapter by considering some of Wittgenstein's remarks on the question of whether a machine can think, and we noted that he was interested in "conceptual" questions, rather than "empirical" questions. We then looked at the introduction of Turing's famous paper on machine intelligence, where we find him complaining (in effect) that if meaning is use, then the question of whether a machine can think ought to be settled by way of a Gallup poll, which is absurd. We argued that pace Turing, it is not absurd to maintain that empirical work directed at how people use terms like “machine” and “think” is relevant to settling the question of whether machines can think, and we suggested several ways that empirical work—including Gallup-style surveys—might engage with the question. We then turned to a disagreement between Turing and Wittgenstein regarding whether mathematics proceeds by invention or discovery. We observed Turing putting forward the same argument (or a very similar one) against Wittgenstein’s view that advances in mathematics amount to giving new meanings to words. We suggested that the disagreement between Turing and Wittgenstein reflected a difference in their fundamental aims. We then tried to sketch the basic outline of their different philosophical projects.

In conclusion, we want to briefly come back to experimental philosophy and indicate how Turing and Wittgenstein’s projects—as we understand them—both have natural empirical complements. We’ll begin with the latter. We’ve indicated that Wittgenstein was intent on discovering the sources of our philosophical problems with an eye toward enabling us to set them aside. And we believe that this project can be aided by empirical work. In this we are motivated by Wittgenstein’s (1958, §66) exhortation, “don’t think, but look!” And we see this as demanding the use of modern empirical methods if we are to look both *carefully* and *systematically*. With regard to Wittgenstein’s philosophical goals, we hold that empirical methods are essential to carefully and systematically describing the ordinary uses of terms, including how their use varies across contexts and across different people. Such descriptive work can help us to better see the lay of the land, including how our language helps shape our philosophical assumptions. It can help us to uncover places where we lay undue stress on certain uses, failing to notice the work that the terms are doing for us and the multitude of uses these terms are put to on which the philosophical puzzles bothering us do not arise. And they can help us to recognize when our usage is idiosyncratic or spurred on by our philosophical inclinations and training, perhaps through an extended feedback loop that escaped our notice. By careful empirical work, we drag our implicit (often metaphysical) assumptions out into the open where they may be criticized.[24]

Turing's project, by contrast, aims to make progress by substitution: replacing intractable problems with more tractable ones. But Turing's project invites skepticism about whether the replacement problems are similar enough to the original. For if the new problem is too different from the old one, then solving the new problem will be unsatisfying. As Wittgenstein quipped (1939, 66): "If [inventing what 'analogous' is to mean] were the problem, we could settle it much easier by making 'analogous' mean 'cushion.'" Empirical work is needed in order to describe the use of the original terms, the aims that they are meant to satisfy, the conceptual roles that they fill, and so on. Further empirical work is needed in order to determine whether a replacement succeeds or fails at its work. Such an empirical project bears close resemblance to experimental explication (see Schupbach, 2015). In explication, one invents new meanings by formalizing old terms. One wants the cleaned-up term to be *close* to the old, fuzzy term. One wants the new meaning—the clarified, formalized thing—to be relevantly similar to the old meaning. At least such that one is not simply changing the subject by using the new term. In

[24] Livengood and Machery (2007) explicitly adopt this goal with respect to the folk metaphysics of causation. When the sources of our metaphysical assumptions appear to be problematic, the Wittgensteinian experimental project turns negative by providing a debunking argument against our ordinary commitments. See Rose and Schaffer (2015) and Korman and Carmichael (forthcoming) for an illustration and discussion of a debunking challenge along these lines.

experimental explication, one calls on empirical methods to help determine the old meaning and characterize how similar it is to the new meaning.

Both of the experimental projects inspired by Turing and Wittgenstein are interesting and important. And the two may be fruitfully conducted together: exploring the use of our language with regard to philosophical questions both as a prophylactic for metaphysical excess and as a guide to clarifying questions by substitution. That the two projects may be fruitfully combined is perhaps unsurprising given that Turing and Wittgenstein shared an affinity for common sense (for which, see Floyd, 2016) and for pragmatism. The experimental turn here would pay attention to the aims and purposes of our language, to its practical significance, and to the sources of our philosophical commitments and puzzlements. In the common ground worked up by Turing and Wittgenstein, experimental philosophers will find a bountiful harvest.

References

Alexander, Joshua (2012). *Experimental Philosophy: An Introduction*. Cambridge: Polity Press.

Alexander, Joshua (2016). "Philosophical Expertise," in J. Sytsma and W. Buckwalter (eds.), *A Companion to Experimental Philosophy*, pp. 557–567. Oxford: Blackwell.

Buckwalter, Wesley and Justin Sytsma (2016). "Introduction," in J. Sytsma and W. Buckwalter (eds.), *A Companion to Experimental Philosophy*, pp. 1–2. Oxford: Blackwell.

Floyd, Juliet (2016). "Turing on 'Common Sense': Cambridge Resonances." In J. Floyd and A. Bokulich (eds.), *Philosophical Explorations of the Legacy of Alan Turing—Turing 100*.

Garver, Newton (1996). "Philosophy as Grammar," in H. Sluga and D. Stern (eds.), *The Cambridge Companion to Wittgenstein*, pp. 139–170. Cambridge: Cambridge University Press.

Kim, Hyo-eun, Nina Poth, Kevin Reuter, and Justin Sytsma (2016). "Where is your pain? A Cross-cultural Comparison of the Concept of Pain in Americans and South Koreans," *Studia Philosophica Estonica*, 9(1): 136–169.

Knobe, Joshua and Shaun Nichols (2008). *Experimental Philosophy*. New York: Oxford University Press.

Knobe, Joshua and Shaun Nichols (2013). *Experimental Philosophy, Volume 2*. New York: Oxford University Press.

Korman, Dan and Chad Carmichael (forthcoming). “What Do the Folk Think about Composition and Does it Matter?” In D. Rose (ed.) *Experimental Metaphysics*.

Livengood, Jonathan and Edouard Machery (2007). “The Folk Probably Don’t Think What You Think They Think: Experiments on Causation by Absence.” *Midwest Studies in Philosophy* 31, 107-127.

Machery, Edouard and Elizabeth O’Neill (eds.), *Current Controversies in Experimental Philosophy*, New York: Routledge.

O’Neill, Elizabeth and Edouard Machery (2014). “Experimental Philosophy: What is it good for?” in *Current Controversies in Experimental Philosophy*, E. Machery and E. O’Neill (eds.), New York: Routledge.

Peirce, Charles Sanders (1878). “How to Make Our Ideas Clear,” *Popular Science Monthly*, 12(18): 286–302.

Reuter, Kevin, Dustin Phillips, and Justin Sytsma (2014). “Hallucinating Pain,” in *Advances in Experimental Philosophy of Mind*, J. Sytsma (ed.), London: Bloomsbury.

Rose, David and David Danks (2013). “In Defense of a Broad Conception of Experimental Philosophy.” Metaphilosophy, 44(4): 512-532.

Rose, David and Jonathan Schaffer (2015). “Folk Mereology is Teleological.” *Noûs*.

Schupbach, Jonah (forthcoming). “Experimental Explication,” *Philosophy and Phenomenological Research*.

Sluga, Hans (1996). “Ludwig Wittgenstein: Life and Work, An Introduction,” in H. Sluga and D. Stern (eds.), *The Cambridge Companion to Wittgenstein*, pp. 1–33. Cambridge: Cambridge University Press.

Suzuki, Jeff (2008). “A Brief History of Impossibility,” *Mathematics Magazine*, 81(1): 27–38.

Sytsma, Justin (2010a). “Dennett’s Theory of the Folk Theory of Consciousness,” *Journal of Consciousness Studies*, 17(3–4): 107–130.

Sytsma, Justin (2010b). “The Proper Province of Philosophy: Conceptual Analysis and Empirical Investigation,” *Review of Philosophy and Psychology*, 1(3): 427–445.

Sytsma, Justin (2012). “Revisiting the Valence Account,” *Philosophical Topics*, 40(2): 179–198.

Sytsma, Justin (2014). “Attributions of Consciousness,” *WIREs Cognitive Science*, 5: 635–648.

Sytsma, Justin (forthcoming). “Two Origin Stories for Experimental Philosophy,” *teorema*.

Sytsma, Justin and Wesley Buckwalter (2016). *A Companion to Experimental Philosophy*, Oxford: Blackwell.

Sytsma, Justin and Jonathan Livengood (2015). *The Theory and Practice of Experimental Philosophy*. Broadview.

Sytsma, Justin and Edouard Machery (2010). “Two Conceptions of Subjective Experience,” *Philosophical Studies*, 151(2): 299–327.

Sytsma, Justin and Edouard Machery (2012). “On the Relevance of Folk Intuitions: A Reply to Talbot,” *Consciousness and Cognition*, 21(2): 654–660.

Sytsma, Justin and Edouard Machery (2013). “Experimental Philosophy,” in B. Kaldis (ed.), *Encyclopedia of Philosophy and the Social Sciences*, pp. 318-320. Thousand Oaks, CA: SAGE.

Sytsma, Justin and Kevin Reuter (forthcoming). “Experimental Philosophy of Pain,” *Journal of Indian Council of Philosophical Research*.

Turing, Alan (1950). “Computing Machinery and Intelligence,” *Mind*, 59(236): 433–460.

Weinberg, Jonathan (2015). “Humans as instruments: or, the inevitability of experimental philosophy.” In E. Fischer and J. Collins (eds.), *Experimental Philosophy, Rationalism, and Naturalism: Rethinking Philosophical Method*, New York: Routledge.

Weisstein, Eric (2017). “Heptadecagon,” from *MathWorld*—A Wolfram Web Resource. <URL= http://mathworld.wolfram.com/Heptadecagon.html>

Williamson, Timothy (2016). “Philosophical Criticisms of Experimental Philosophy,” in J. Sytsma and W. Buckwalter (eds.), *A Companion to Experimental Philosophy*, pp. 22–36. Oxford: Blackwell.

Wittgenstein, Ludwig (1958). *The Blue and Brown Books*. Oxford: Basil Blackwell.

Wittgenstein, Ludwig (1958). *Philosophical Grammar*. Edited by R. Rhees, translated by A. Kenny. Oxford: Basic Blackwell.

Wittgenstein, Ludwig (1958). *Philosophical Investigations, Third Edition*. Translated by G.E.M. Anscombe. Upper Saddle River, NJ: Prentice Hall.

Wittgenstein, Ludwig (1976). *Lectures on the Foundations of Mathematics*. Edited by C. Diamond. Ithaca: Cornell University Press.

Wittgenstein, Ludwig (1980). *Culture and Value*. Edited by G. H. Von Wright with H. Nyman, translated by P. Winch. Chicago: University of Chicago Press.