

A Novel Exercise for Teaching the Philosophy of Science

Gary Hardcastle and Matthew H. Slater*†

We describe a simple, flexible exercise that can be implemented in the philosophy of science classroom: students are asked to determine the contents of a closed container without opening it. This exercise has revealed itself as a useful platform from which to examine a wide range of issues in the philosophy of science and may, we suggest, even help us think about improving the public understanding of science.

1. Introduction. On the first day of our undergraduate philosophy of science classes, our students are handed a syllabus typical of syllabi for such courses as they have been offered in English-speaking universities over the past thirty years. But they are also presented with something unusual—not just for a philosophy of science course, but for any course. “This,” we tell our students, holding up a sealed cardboard box for all to see, “is the box.” We shift and shake the box slightly, making it apparent that the box has contents, of distinct shapes, sizes, and masses. “During the semester you will complete what we will call the Box Project,” we explain, “your grade for which will account for 20% of your final grade in this course. In a few minutes we will form groups. For the rest of the semester, each group is tasked with answering one question, and one question only: What is in the box? And, there is just one rule to obey as you try to answer this question: you may not open the box. One question, one rule. Are these both clear?”

We pause to let the students affirm that the task and the stricture are clear. We have two more things to tell them, we add: “The groups are in compe-

*To contact the authors, please write to: Gary Hardcastle, Department of Philosophy, Bloomsburg University of Pennsylvania, 400 E. Second Street, Bloomsburg, PA 17815; e-mail: ghardcas@bloomu.edu. Matthew H. Slater, Department of Philosophy, Bucknell University, One Dent Drive, Lewisburg, PA 17837; e-mail: matthew.slater@bucknell.edu.

†We thank the participants of our PSA 2012 Symposium and an anonymous referee for this journal for comments.

tition, and one group will win. Each member of the winning group will receive a \$50 certificate good for dinner at Bob's Land o' Sushi" (or functionally equivalent local eatery, university store, etc.). "Second," we say, "each group will submit two confidential reports on its progress and, at semester's end, present its answer to the question. The group reports and presentations will be given letter grades reflecting quality and thoroughness, and these will determine your group's grade for the Box Project. Each member gets his or her group's grade." "In short," we stress, "your grade for the Box Project will reflect the quality of your work, not whether you win."

We invite questions. Invariably students ask whether they can X-ray the box, shine bright lights on the box, or subject the box to some other treatment, and invariably they are told that they can do anything to the box, so long as they do not open it. Having determined a suitable place to store the box (which is available by "sign-out" to any group) and set a proximate deadline for the first group reports, we form groups and move on to the next topic.

At this point, the students know much, but not all, that there is to know about the Box Project. It is true that there is just one question to answer and just one rule to follow. They do not know what is in the box, of course, but neither do we, having had the box filled and sealed so that we could not know what was in it. They often think they know how the competition will unfold. Surely, they presume, at semester's end the box will be opened, and the group with the most accurate answer—namely, the answer that most closely matches the finally revealed contents of the opened box—will win. That sounds to them exactly like something that would happen in a sort of "philosophical science class" (which is often what they have imagined the philosophy of science must be) and indeed does happen in science, as they understand it: one's guesses about reality can be eventually checked against reality.

The students are correct that the Box Project is an attempt to get them to view science from the point of view of a scientist, but they are wrong that the exercise will end with an opened box and accuracy scores for each group's answer. No scientist has ever had the analogous pleasure of comparing her theory about some part of the world with that part of the world itself, and no student can have that pleasure if she is to get a sense of science from within. So, at the semester's end (and beyond), the box remains sealed. Unless Box Project veterans have tipped off our students, the news that the box will not be opened is a shock, met with disbelief and disappointment. But this gives way to an appreciation for the point it makes, namely, that scientists never get to find out how some bit of theory "really did." Our claims about the world inevitably extend beyond the immediately observable, and for this we pay a certain epistemic price.

2. The Main Pedagogical Challenge of Teaching Philosophy of Science.

The Box Project was deployed first in fall 2002; we have used it (separately) in nine upper-level undergraduate philosophy of science classes since then, each with 8–30 students (most not philosophy majors), meeting 3 hours a week over a 15-week semester. The exercise was conceived to address the pedagogical challenge posed by students in philosophy of science classes who have little or no exposure to the practice of science. These students feel (with some justification) unqualified to assess any theory or account of science, and indeed they often have not yet learned that there is such a thing as a “theory of science.” Science majors are often among these students, especially if their lab experiences have consisted of highly scripted exercises, designed to develop specific lab skills rather than the problem-solving and inquiry tools needed to produce new knowledge. The Box Project is meant to give these students a sense of what it is like to be a scientist—an incomplete sense to be sure, but one sufficient to provide the capacity to understand and evaluate a claim about science as a student might find it expressed by Hempel, Popper, Kuhn, or any philosopher of science.

In our judgment, the Box Project succeeds on this score: it addresses the problem of the philosophy of science student unexposed to and detached from science. We believe that students who complete the Box Project are, on the whole, more engaged and confident in the course, and that as a result they learn more philosophy of science and better appreciate the discipline itself. Our evidence for this might, of course, be found in course evaluations. These were collected for several classes using the Box Project but are unavailable to us in this context on pain of making unwitting (and thus non-consenting) research subjects of our past students. At any rate, such evaluations would at best provide anecdotal evidence for our claim in the form of written comments provided by students. Our own confidence that the Box Project works as envisaged is founded on a store of experiences and anecdotes rather than any systematic or rigorous study and is perhaps the worse off for it. In the absence of rigorous assessment we offer the discussion to follow—much of which turns on the surprising ways the Box Project naturally illuminates a number of science’s most interesting features—as mounting a plausibility argument for the exercise’s worth in a philosophy of science classroom.

A few of the Box Project’s practical benefits deserve comment. The Box Project is easy to implement in the classroom. It requires little in the way of materials to create, and its sole imperative and sole rule are easily conveyed to students. It is flexible, as it can be adapted to courses in the philosophy of science with remarkably different emphases. And largely because of its simplicity, the Box Project is robust: there is surprisingly little that can go wrong in its implementation. The worst possible disaster would be, of course, to have the box opened, whether through intent or mishap. This has

happened to neither of us, though. Perhaps this is because our students are informed that there is no “Plan B” should the box be opened, and therefore that if the box were opened (for any reason) the project (which accounts for 20% of their final grade) may go ungraded, a scenario apparently so terrifying to students that they actually comply. As we noted, we have never had to abandon the project because of a clearly opened box.

What about the opposite pedagogical problem, a box that gets little or no attention? We find that while the grade incentive and financial reward are sufficient (and in cases necessary) to spark efforts to figure out what is in the box, once students have engaged that question, little more is required to motivate them; many students embrace the Box Project as a puzzle, for its own sake, even after the course ends. And, more generally, we have found that what might at first glance appear as a crisis in an implementation of the Box Project will in fact set the stage for a genuine debate over some issue at the heart of the philosophy of science. For example, when one of our groups revealed in its final presentation that it had passed a fiber-optic cable through a slight (already-existing) gap in the box, enabling color photographs of its contents, the result was not a collapse of the project but rather a fierce debate between the students over whether the rule guiding the science at hand—do not open the box—had been violated. What had appeared at first to be a headache for the instructor was in fact a chance for students to experience methodological debate. In fact, having conveyed the task and the rule of the Box Project, there was little for an instructor to do but listen to the debate with pleasure.¹ Witnessing many such self-averting crises over several implementations of the Box Project has left us confident not only in its robustness but in the project’s value as a platform for teaching the philosophy of science, a topic to which we now turn.

3. The Box Project: Additional Pedagogical Benefits. The Box Project has revealed to us a class of additional benefits specific to the teaching of the philosophy of science. In broad terms, the Box Project has proven itself a versatile tool that an instructor can use to illustrate a number of central issues in the philosophy of science in a novel and compelling way. To argue this claim, we need not (nor have we the space to) describe the entire range of topics we have seen illustrated via the Box Project; a few salient examples and some reflection will, we hope, make our case. Note that we intend these only as illustrations, not as essential lessons of sustained reflection on a closed cardboard box. The Box Project serves only as a neutral platform

1. Not that there have not been close calls. In 2004 a group took the box to a US airport, planning to have it X-rayed. They told security that they had no idea what was in the box, and that they were under strict orders not to open it. Amazingly, and fortunately, the box and the students survived the episode neither harmed nor incarcerated.

from which a wide range of issues in the philosophy of science can be broached and examined.

3.1. Realism and Antirealism. Our earlier description of the Box Project did not mention how the winning group is chosen. We adopt different strategies for this part of the exercise. In both implementations, each group gives a final presentation on or near the last day of class. As many as 20 minutes are allotted for each presentation, including questions from other groups and any guests joining the class. In Slater's class, those guests include other professors (mainly from the sciences), invited to constitute a panel of judges. Immediately after the presentations, those judges confer in private to choose a winning group; their decision is final. In Hardcastle's implementation, just before the presentations begin, each group elects a representative to an executive committee, which meets immediately after the presentations to choose a winning group; their decision is also final.

We have described how the winning group is selected not just to complete our description of the Box Project but because it is in these later stages of the exercise that students are most often confronted with alternative, competing interpretations of their data. Something along these lines will have occurred throughout the semester, of course, as individuals within groups debate or challenge the group's developing answer to the question of what is in the box. But, in our experience it is during the group presentations, when groups face questions from other students and perhaps other faculty, that alternative interpretations gain their full force. By this point, groups have often managed to X-ray the box, and it is often in the presentation of these X-rays that we witness that force. The following interaction is representative. Team A had presented their X-ray of the box and identified a small circular shadow as a Life Saver™ candy, when a representative of Team B interjected:

TEAM B. How do you know that's a Life Saver™? Couldn't it be some kind of washer?

TEAM A. That can't be a washer. If it was metal, it would be much brighter on the X-ray. The only thing it could be is a Life Saver™.

TEAM B. We didn't say it was a metal washer; maybe it's silicone. That's consistent, we believe, with the density suggested by the X-ray.

TEAM A. Wait. Aren't all washers metal?

TEAM B. Um . . . no, I don't think so. I'm pretty sure that some are rubber or plastic.

We have seen this sort of scenario repeated many times. In another case, one group confidently identified an American quarter and dime in an X-ray, demonstrating that the diameters of the relevant shadows matched the

diameters of an American quarter and dime. The demonstration was compelling until another student pointed out that coins of other countries may well have the same diameters, and yet another noted that the X-ray device may have been set to either enlarge or shrink images.

As distressing as such moments are for students, they should delight the instructor. Alternative interpretations of data are an entry point into a web of issues in the philosophy of science and, despite their centrality, are difficult to teach. Richly drawn historical examples of alternative interpretations of data and competing theories are fascinating and should, at any rate, be part of any philosophy of science course. But too often they fail to convey the urgency such alternatives had at the time and, therefore, the motivation for careful thought about how exactly one should proceed in the face of such alternatives, as well as how one should think about one's commitment to one interpretation over another. Knowing that the earth orbits the sun, that oxygen is consumed in combustion, and that one's genetic information is stored (mostly) in one's DNA, students are apt to regard the Tychonic system, phlogiston theory, and protein-centered genetics as mere historical footnotes. Contemporary examples of competing interpretations of data remedy this problem to some extent if they can be made accessible to students, but even in the best cases students will not regard the debate as their problem (as indeed it is not). When students are invested in what is in the box, though, interpretations they had not considered are received as challenges that demand a response.

In this setting students are inclined to defend their interpretation in a number of ways, many of which illustrate appeals to what are sometimes referred to as a theory's "extra-empirical" virtues. A group might cede that their evidence is equivocal between various interpretations and appeal to their preferred answer's simplicity, its greater coherence with background theories, and so on, as a reason to accept it over alternative answers—an obvious opportunity for connection to themes discussed in many courses.

Recall that during the presentations students often presume that any doubts about a group's answer will be laid to rest (or, alternatively, confirmed) shortly, when the box is opened. The revelation that the box will never be opened gives students a valuable insight into the nature and appeal of antirealism about science, a classification we employ to describe accounts that counsel against conceiving the aim of science as discovering the truth about the natural world, no matter how well confirmed our theories seem. Our students have a natural inclination to realism and are typically perplexed when they encounter antirealist accounts of science. Confronted by a range of heretofore-unanticipated alternative theories of what is in the box, all of them compatible with the evidence at hand, plus the realization that there are likely more alternative accounts of their data yet to be conceived (Stanford 2006), plus the fact that the box may never be opened, students rec-

ognize and appreciate the epistemic commitment realism entails and the challenge involved in defending that commitment.

3.2. *The Box as a Kuhnian Puzzle.* In section 2 we alluded to the appeal some students have for the box as a puzzle to be solved, sometimes coming to give less regard to their grade or chance of winning than to simply figuring out what is in the box. It is easy to empathize. The task seems easy—after all, the contents of the box are in one’s hands—and yet the path to specifying those contents precisely, without opening the box, is not at all clear. It is a recipe for obsession, and some students find themselves obsessed.

If the course covers Kuhn’s *The Structure of Scientific Revolutions* (1962/1970; especially its fourth chapter), one can at this point demonstrate to the students in their own behavior a personality trait that plays a central role in Kuhn’s account of science. Having characterized normal science as the potentially rather dreary (typically experimental) determination of some constant or other result (one that is, courtesy of the prevailing paradigm, all but known in advance), Kuhn faced the obvious question of why anyone would want to do that, for even an afternoon, let alone a lifetime. Kuhn’s response comes close to defining normal scientists as just those persons who enjoy just such pursuits, that is, pursuits in which the aim is less to discover the goal than to discover a path to that goal. Only such people would find the problems of normal scientific research fascinating. “Though its outcome can be anticipated, often in detail so great that what remains to be known is in itself uninteresting,” Kuhn wrote of the normal science research problem, “the way to achieve that outcome remains very much in doubt. Bringing a normal research problem to a conclusion is achieving the unanticipated in a new way, and it requires the solution of all sorts of complex instrumental, conceptual, and mathematical puzzles. The man who succeeds proves himself an expert puzzle-solver, and the challenge of the puzzle is an important part of what usually drives him on” (1962/1970, 36). Normal science, in Kuhn’s view, is less for seekers of significant truths about nature than for those with a facility for and an attraction to working out how to measure, calculate, or determine some subtle feature of nature, even (and especially) when much about that feature is already known. Students engaged in the Box Project, gripped by the puzzle or not (and some are, always), are better positioned to understand and critique Kuhn’s account of science. Seeing the Box Project as a Kuhnian puzzle (discovering what is in the box is, after all, of no intrinsic interest or importance, much about its contents is already known, and admissible solutions are constrained by a rule), they can ask and answer such questions as “Can the pleasure of puzzle solving really account for a lifetime’s pursuit of normal science?” and “If normal science is largely puzzle solving—the pursuit of clever

ways to discern largely uninteresting and even trivial results—of what value is it to society?”

While *Structure's* chapter on normal science as puzzle solving is its shortest, the points made here loom large in Kuhn's full account. Since the rules that guide the puzzles of a normal science are, in Kuhn's account, just the laws and convictions bound up with its paradigm, to engage a puzzle is to engage a paradigm (1962/1970, 39–44). Further, scientific revolutions are, in Kuhn's view, precipitated by crises, and these in turn are precipitated by unsolved puzzles. Finally, scientific progress is measured, according to Kuhn, in the accumulation not of truths but of “puzzle-solutions” (1962/1970, 160–73). Thus, in just this single facet of the Box Project—the fact that some students embrace the Box Project as a puzzle—we find a platform from which to engage the main pillars of Kuhn's account of science.

3.3. The Box and Science Funding. Very early in the Box Project, as groups learn (through offhand discussion or through informal midsemester group presentations) what their peers are up to, it becomes clear to all that resources are not distributed fairly. Whether groups are chosen randomly by a count-off procedure or allowed to self-select, one or more groups will have resources the others lack, such as a relative or close friend with access to special equipment, most often an X-ray or sonogram machine and, crucially, a trained professional willing to help interpret the resulting images. Groups will also have differential access to campus resources, including, for example, fiber-optic cameras that might be passed through an opening in the box or faculty willing to advise one group over another. Students who note the disparity do not always complain about it, but those who do can be reminded that the Box Project's one rule—do not open the box—has not been violated, so the exercise will continue. All students are reassured that their grade for the Box Project is independent of whether one's group wins or loses; losing groups can (and have) done excellently, and (though less often) winning groups have received poor grades.

Here one can profitably introduce questions about science funding and more broadly about the values that various funding models reflect. One might ask the students how one ought to distribute some large sum (say, \$10,000) to the class's four or five groups, given that one's sole aim was to determine what is in the box. Should the full sum be invested in whatever group is most able, best resourced (intellectually and materially), and most likely to get it right? Why not hedge one's bet by giving some money to one of the long-shot groups? If so, how much? Does fairness enter into any of these calculations? Students—and especially those riled by the unfair distribution of valuable resources—thus grasp the complexities of science funding and the broader questions concerning the intersection of our science and our values.

4. Challenges and Variations. Though our experience with the Box Project has been very positive, its implementation does present some challenges. As mentioned, the project's flexibility and robustness facilitate easy handling of many of these. For example, in large classes "box time" is hard to come by. Though this presents an opportunity for discussion—access to time with certain scientific instruments (e.g., the *Hubble Space Telescope*) is hard to come by too—it can also lead to unproductive frustration. Another challenge devolves from some students' perception that the Box Project is "busywork" and thus to be deprioritized. While this appeared to be a minority response, it does highlight a difficult balancing act instructors have to play in this context. On the one hand, we want to be able to offer students a compelling rationale for every assignment (especially those that are apt to seem like busywork); on the other hand, doing so conflicts with the strategic vagueness that allows students to guide and motivate themselves. It seems to us that the best way of minimizing this response is to structure the project in such a way that students are immersed in it as soon as possible in hopes that early involvement will convey the value of the exercise. Many are also motivated by the competition. The team aspect can promote this competitive investment (thus, we like having the students work in their teams on small-group activities in class, even those at right angles to the Box Project).

A third challenge involves protecting the secrecy of the denouement. Surprise is one of the more important features of the project, both promoting student "buy-in" along the way (since students presume that the box will be opened in a dramatic unveiling ceremony) and driving home important insights about science at its conclusion. Much would arguably be lost if students showed up knowing the full nature of the project. One way of dealing with this problem would be to use the Box Project sporadically in order to disrupt the standard processes of "campus-memory" formation. A second strategy is to get the students involved in the project in such a way that they agree not to divulge the secret. Some of Slater's students asked whether they could fill the box for the next class—no doubt with yet more puzzling objects. Some even suggested that the box itself be but one component of a larger subsequent box.

These brief comments hint at the many variations on the basic Box Project as we have described it. The project could, for example, last a whole term (as it does in our implementations) or span a shorter period. Groups could be assigned, or students could be left to their own devices to form groups. Instructors might model the behavior of funding agencies, possibly even making access to the box contingent on a compelling presentation in front of their peers. As mentioned above, Hardcastle has final presentations judged by a student executive committee, while Slater invites other professors to serve as a panel of judges. Students are reminded that little things—such as stage presence, grooming, and professional attire—can subtly affect the pan-

el's judgment. They get nervous—especially as the panel and other teams begin to question their inferences, methods, and assumptions.

We can easily imagine yet other variations. Perhaps students could be allowed to change groups. Or there could be multiple boxes, some of whose contents are easier than others to discern.² Unfair? Yes, but of course not all scientific puzzles are equally easily solved. Or perhaps the professor could take on the role of a malevolent deity, swapping out the box with a (largely) indistinguishable box with different contents at some point during the term. This would have the tendency of producing anomalies—either within groups whose exposure to the box straddled the swap or during the final presentations (assuming that at least one team got to examine a different box than the others). Would students question their paradigm of a hands-off (and truthful) professor? Maybe. But that itself would be an interesting scenario to discuss.

5. The Box Project: A Public Benefit? We have made what we hope is at least a compelling *prima facie* case for the place of the Box Project in the philosophy of science classroom. But the Box Project's capacity to bring students to central questions and issues in the philosophy of science suggests that the exercise is relevant to a more serious problem: the poor grasp of science among the nonscientific public.

It is hard to deny that there is such a problem. In its *Science and Engineering Indicators 2012*, the National Science Foundation (NSF) reports that “many Americans continue to give multiple incorrect answers to questions about basic factual knowledge of science” (2012, 7.24). Americans responding in 2010 to a set of mostly true/false questions “were, on average, able to correctly answer 5.6 out of the 9 items.” Only 73% of the 1,932 subjects knew that the earth went around the sun, for example, and only 51% responded that electrons are smaller than atoms. Subjects who had had nine or more science and mathematics classes managed to answer only 83% of these factual questions correctly, on average. And this dismal state of affairs has not changed in decades (7.19).

The American public's understanding of the generation and assessment of scientific evidence, rather than particular scientific facts, is no less alarming. The NSF reports that 42% of 1,454 respondents “understood scientific inquiry,” a designation subjects earned for correctly answering two questions concerning probability and providing acceptable answers to at least one of two open-ended questions about scientific methodology or experiment. Among subjects with no more than a high school education, only 26% were gauged to understand scientific inquiry, while 71% of subjects who had completed three college-level science and math courses exhibited

2. This suggestion was offered by Janet Stemwedel.

such understanding. For those who had completed a graduate or professional degree, that proportion leapt to 73% (National Science Foundation 2012, 7.25).

Such widespread ignorance of basic scientific facts and of the nature of science poses an especial threat in the United States, where budgets supporting a significant portion of the world's scientific research, as well as policies and initiatives that affect the entire world, risk being shaped by an electorate one-fourth composed of people who believe that the sun orbits the earth (cue "sad trombone"). And as the NSF's report suggests, this ignorance of basic scientific facts rests in part on an inability to evaluate scientific evidence and, more generally, an ignorance of the very nature of science. Given that the philosophy of science has made significant and lasting contributions to, among other things, our understanding of how and when evidence speaks for (or against) a scientific theory, what we are entitled to believe about ourselves and the world on the basis of well-confirmed theories, and how our efforts to learn about the natural world interact with the political and social forces that shape our lives, it seems clear that the melioration of the public's sorry understanding of science falls within the capability and purview of philosophers of science.

Could the Box Project help us think about how to do this? To the extent that that exercise gives our students (who are, after all, members of the nonscientific public) a look at science from which, guided by philosophers of science, they may come to a deeper understanding of scientific evidence, agreement, change, and progress, we suggest that the Box Project and similar exercises may be used to improve both the general public's recognition and understanding of science and its grasp of the facts about our world we take science to have established. Our suggestion does not require developing nonclassroom versions of the Box Project; students are, as we noted, members of the public, and the Box Project can be easily adapted to and implemented within any number of other classes, from the earliest grades up. Of course, while the Box Project is robust, requiring little from an instructor to set up, it does not teach the philosophy of science by itself. Outside a philosophy of science classroom, the capacity of the Box Project to convey its lessons about science and the philosophy of science depends on having a philosopher of science on hand to point those lessons out.

Our suggestion that philosophers of science have a role to play in ameliorating widespread ignorance of science betrays an optimism about the discipline and its achievements that may be somewhat out of fashion. It is more common, even among philosophers of science, to note the gap between the large number of arcane questions taken up by the highly technical and complex contributions to the field's leading journals, on the one hand, and the very different questions that occupy the general public (including scientists), on the other hand, and to take that gap as evidence that phi-

osophy of science has lost all contact with broader concerns and become, in a phrase of John Dewey's, a "sentimental indulgence for a few" (1997, 328). Such is the diagnosis that Philip Kitcher (2011) offers of (primarily Anglophone) philosophy's "core areas"—metaphysics, epistemology, philosophy of language, and philosophy of mind—whose proponents, Kitcher claims, have managed to replace the pursuit of questions of great and compelling human interest with the pursuit of an array of stunningly arcane questions, posed in forbidding formalisms, utterly disconnected from anything of genuine human interest, and unanswerable to boot. What is Kitcher's outlook for these core areas of philosophy? "As the philosophical questions diminish in size," Kitcher writes, "disagreement and controversy persist, new distinctions are drawn, and yet tinier issues are generated. Decomposition continues downwards, until the interested community becomes too exhausted, too small, or too tired to play the game any further (2011, 251). Barring intervention, then, today's core areas of philosophy will shrink to nonexistence. Kitcher, following Dewey, thus calls for a reconstruction of philosophy, replacing its doomed core areas with those Kitcher takes to be at present peripheral, namely, those that respond to "the felt needs of individual people to make sense of the world and their place in it" (254).

The extent to which Kitcher takes the philosophy of science as part of philosophy's current core, and thus subject to his indictments, is not clear. Perhaps he does not include it; the discipline is never mentioned by name, and philosophy reformulated along Kitcherian lines would retain a place for the pursuit of natural knowledge (and the examination of that pursuit) along what Kitcher calls philosophy's "knowledge axis." On the other hand, much prized current work in the philosophy of science bears more than a passing resemblance to those sentimental indulgences for the few, taking up questions of extremely local interest and employing highly specialized mathematical and philosophical apparatuses to approach them. No matter though, as we are in broad agreement with Kitcher's aim to make philosophical work bear on real human needs, and we would offer the Box Project as a humble contribution toward it. Philosophy reformulated on Kitcher's suggestions remains rightly concerned with furthering human efforts to learn about the world. The expression of that concern within philosophy takes various forms, from directly seeking it, to serving as a "midwife" for nascent fields of natural science, to untangling conceptual confusions found in other, mature sciences. We should expect these and other forms to wax and wane over time were philosophy to refocus itself on genuine and pressing human concerns. But to the extent that the Box Project and similar exercises share the elements most basic to inquiry and can thus be made to stand for natural inquiry in very general form, the exercise offers an effective and novel means of addressing those real and pressing concerns. And, as we saw above, the Box Project can be used to bridge the gap be-

tween the highly technical and seemingly arcane language and tools of contemporary philosophy of science by placing students in a context in which they feel the grip of the broad general questions that ultimately give rise to the specialized language and techniques. Until one is confronted by alternative interpretations of an X-ray of the box, for instance, consistent with all the available data, it is hard to appreciate the problem posed by such alternatives; once confronted, the questions surrounding underdetermination, alternative theories, and realism can be appreciated, and the best of the highly technical literature produced in their wake understood as a sustained collective effort to work out how science ought to proceed.

REFERENCES

- Dewey, John. 1997. *Democracy and Education*. New York: Free Press.
- Kitcher, Philip. 2011. "Philosophy Inside Out." *Metaphilosophy* 42 (3): 248–60.
- Kuhn, Thomas. 1962/1970. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- National Science Foundation. 2012. "Science and Technology: Public Attitudes and Understanding." In *Science and Engineering Indicators 2012*, chap. 7. Arlington, VA: National Science Foundation. <http://www.nsf.gov/statistics/seind12/pdf/c07.pdf>.
- Stanford, P. Kyle. 2006. *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. New York: Oxford University Press.