

Polling as Pedagogy: Experimental Philosophy as a Valuable Tool for Teaching Philosophy

THOMAS NADELHOFFER

Dickinson College

EDDY NAHMIAS

Georgia State University

Abstract: First, we briefly familiarize the reader with the emerging field of “experimental philosophy,” in which philosophers use empirical methods, rather than armchair speculation, to ascertain laypersons’ intuitions about philosophical issues. Second, we discuss how the surveys used by experimental philosophers can serve as valuable pedagogical tools for *teaching philosophy*—independently of whether one believes surveying laypersons is an illuminating approach to *doing philosophy*. Giving students surveys that contain questions and thought experiments from philosophical debates gets them to actively engage with the material and paves the way for more fruitful and impassioned classroom discussion. We offer some suggestions for how to use surveys in the classroom and provide an appendix that contains some examples of scenarios teachers could use in their courses.

When teaching philosophy, we often begin discussions by asking our students some question, like, “Do you think we have free will if God already knows everything we will choose?” or by describing thought experiments, like zombies (creatures physically identical to us with no conscious experiences) and asking if they are possible or impossible. We may then ask for a show of hands, perhaps calling on a few students to explain why they “voted” as they did in order to get the philosophical discussion started. We informally poll our students in this way in order to get them to think about the issues, to demonstrate the conflicting intuitions that motivate competing philosophical positions, and to get them interested and engaged in the philosophical debate. In this article, we suggest that this pedagogical technique of polling our students is a valuable one—so valuable that it should be used more often and in a more rigorous way. Let us explain.

There is growing interest among philosophers working under the rubric of “experimental philosophy” in the intuitions and beliefs of laypersons. Researchers working in this nascent field have begun examining “folk” intuitions and conceptual usage concerning a myriad of philosophical problems in areas as diverse as epistemology (Nichols, Stich, and Weinberg 2003), ethics (Nichols 2004a, 2004c), free will (Nahmias *et al.* 2005, 2006; Nichols 2004b; Nichols and Knobe 2007), the philosophy of language (Machery *et al.* 2004), action theory (Knobe 2003, 2004; Nadelhoffer 2004, 2005), and the philosophy of law (Nadelhoffer 2006). One of the goals of experimental philosophy is to use controlled and systematic methods rather than armchair speculation to determine where people’s ordinary intuitions about philosophical problems lie—if they ordinarily lie anywhere at all.¹ After all, philosophers frequently appeal to common sense, intuitions, and ordinary conceptual usage. Determining precisely what the facts are about these things ought to be at least in part an empirical investigation.

In this paper, we first briefly familiarize the reader with some of the interesting and surprising findings of experimental philosophy—focusing primarily on recent research on folk intuitions concerning free will and moral responsibility. Second, we discuss how the kind of surveys used by experimental philosophers can serve as a valuable pedagogical tool for *teaching philosophy*—independently of whether one believes surveying laypersons is an illuminating approach to *doing philosophy*. Giving students surveys that contain questions and thought experiments from philosophical debates gets them to actively engage with the material and paves the way for more fruitful and impassioned classroom discussion. We also offer some suggestions concerning how to use surveys in the classroom and we provide an appendix that contains some additional examples of scenarios that teachers could use in the classroom.²

Conceptual Analysis and Folk Intuitions

Ever since Socrates began challenging the beliefs and assumptions of his fellow Athenians—often much to their chagrin—philosophers have been in the business of analyzing concepts. One of Socrates’ motivating assumptions was that we should seek to find necessary and jointly sufficient conditions for the correct application of the general terms that we use in ordinary language. This is an assumption that is still operative among many analytic philosophers today. Indeed, for a number of contemporary philosophers, the method of doing philosophy is still broadly Socratic. When engaged in conceptual analysis, philosophers first pick a particular concept to investigate. Then, they put forward

tentative definitions and determine whether these definitions fall prey to counter-examples or involve any hidden internal inconsistencies.

Given this view of conceptual analysis, intuitions about particular cases often take center stage in contemporary philosophy.³ But to the extent that intuitions play a central role in conceptual analysis, it seems reasonable to ask whose intuitions philosophers ought to count.⁴ While some philosophers may refer only to the tutored intuitions of philosophers, many refer instead to the pre-theoretical intuitions of laypersons. On this latter view, ordinary concepts and folk intuitions are the proper subject matter of conceptual analysis, and they are meant to offer *prima facie* support for or evidence against a philosophical theory.

This seems to be the approach to conceptual analysis adopted by Socrates himself. After all, in asking his interlocutors to answer questions such as “What is courage?” or “What is knowledge?” he was trying to force them to give an account of their own ordinary concepts and not some technical or philosophical counterparts. Moreover, the effectiveness of the Socratic method depended on the interlocutors’ own intuitions about the counter-examples that Socrates put forward. In this respect, the philosophical appeal to intuitions has often—although admittedly not always—been fueled by an interest in folk judgments concerning particular cases. Similarly, when we teach philosophy we also use the Socratic method to probe our students’ intuitions and to prompt them to come up with justifications for these intuitions.

Recently, philosophers have begun to recognize the importance of trying to get at folk intuitions in a more rigorous way than simply asking students for an informal show of hands. Working at the crossroads of philosophy and the social sciences, experimental philosophers have begun borrowing methods from psychology to formally probe folk intuitions with the goal of shedding light on philosophical debates that are mired in intuitional conflict. Of course, experimental philosophers do not suggest that discovering what people’s intuitions actually are will *solve* any philosophical problems. Nor do they suggest that ordinary language or pre-theoretical judgments should serve as the final court of appeal in philosophical debates. The goal of experimental philosophers is more modest.

Researchers first and foremost want to get at the relevant folk concepts and intuitions in a controlled and systematic rather than merely speculative manner. Presumably, once we know what the pre-theoretical intuitions about a particular philosophical problem are, philosophers will no longer be able to claim that their position aligns with common sense unless their views empirically merit such support. In the event that a particular analysis of an ordinary concept *does* turn out to settle with folk intuitions that alone would not prove it to be *true*, but it would seem to shift the argumentative burden to those philosophical

theories or analyses that have counterintuitive implications. And if ordinary intuitions are hopelessly inconsistent or subject to irrelevant biases, then philosophers should be wary about using them as evidence for philosophical claims.⁵

For present purposes, we are going to set aside the question of whether experimental philosophy is a fruitful way of *doing philosophy*, since adequately answering this question would take us too far afield.⁶ We are going to focus instead on the value of using systematic surveys for *teaching philosophy*. In our experience, polling students is not only an excellent way of getting them to engage with the material, it is also an excellent way of getting them to actively participate in class discussion. Moreover, by taking the time and energy to learn more about a particular group of students' intuitions concerning philosophical problems, we place ourselves in a better position as teachers to help students gain a perspicuous view of the relevant issues involved.

Folk Intuitions about Free Will and Moral Responsibility


Perhaps the best way to begin would be to consider examples of surveys that experimental philosophers have relied on in their efforts to probe folk intuitions. One area that has received a lot of attention recently is the free will debate. Incompatibilists often claim that their view is intuitive and that compatibilism is counterintuitive; for instance, Robert Kane writes, "In my experience, most ordinary persons start out as natural incompatibilists. . . . Ordinary persons have to be talked out of this natural incompatibilism by the clever arguments of philosophers" (Kane 1999: 217). And incompatibilists sometimes suggest that informal polls of their students provide evidence for this claim. For instance, Derk Pereboom claims that "[b]eginning students typically recoil at the compatibilist response to the problem of moral responsibility" (Pereboom 2001: xvi), and Timothy O'Connor writes, "Does freedom of choice have this implication [that causal determinism is false]? It seems so to the typical undergraduate on first encountering the question" (O'Connor 2000: 4).

As we've said, we support polling students, but it is important to do so in a systematic way, since it is easy to present philosophical problems in ways that will adduce particular answers. Indeed, we have our own unscientific "evidence" indicating that a compatibilist teacher can present the issue so that most students do *not* see a problem with determinism and raise their hands in support of a compatibilist conception of free will. So it is important to present the issues without using cases that would beg significant questions (*e.g.*, by presenting determinism as equivalent to fatalism or manipulation by the forces of nature) and, if done in the classroom, without students knowing the

position held by the teacher. Of course, in the literature, philosophical thought experiments are often designed as rhetorical devices, aiming precisely to pump intuitions in specific ways, rather than as neutral probes of pre-philosophical intuitions. In such cases, teachers can use such intuition pumps to discuss with their students the features of the cases that do the rhetorical work. However, in order to probe students' pre-philosophical intuitions effectively, the cases should contain as few philosophically irrelevant features as possible and be pitched at a level students can understand.

In a series of recent papers, we discuss the results of several studies we ran that examined folk intuitions concerning free will, determinism, and moral responsibility, studies in which we attempted to present the issues in a neutral way.⁷ Participants were students at Florida State University who had not studied the free will debate. They were asked to read a scenario and answer the questions that followed, and they were told there were no right or wrong answers. In one study students received the following vignette:

Imagine that in the next century we discover all the laws of nature, and we build a supercomputer that can deduce from these laws of nature and from the current state of everything in the world exactly what will be happening in the world at any future time. It can look at everything about the way the world is and predict everything about how it will be with 100 percent accuracy. Suppose that such a supercomputer existed, and it looks at the state of the universe at a certain time on March 25, 2150, twenty years before Jeremy Hall is born. The computer then deduces from this information and the laws of nature that Jeremy will definitely rob Fidelity Bank at 6:00 p.m. on January 26, 2195. As always, the supercomputer's prediction is correct; Jeremy robs Fidelity Bank at 6:00 p.m. on January 26, 2195.



Some students were asked whether they believed that in robbing the bank Jeremy acted “of his own free will,” and others were asked instead whether Jeremy was morally responsible for robbing the bank.⁸ The results indicate that a significant majority of students (76 percent) who received the first question judged that Jeremy robs the bank of his own free will, and 83 percent of the students who received the second question judged that Jeremy was blameworthy for robbing the bank. Scenarios that had Jeremy performing positive actions (saving a child from a burning

building) and neutral actions (going jogging) produced similar results. On the surface at least, it appears that people are pre-theoretically less incompatibilist than philosophers have traditionally assumed.

One might object that in the Jeremy cases, we did not make the deterministic nature of the scenario salient enough to the participants. Perhaps they were more focused on the fact that Jeremy's actions were predicted by the supercomputer than the fact that the prediction was made based on deterministic laws. To explore this possibility, we developed another scenario that presents determinism in a different and more salient way by highlighting that the agents' behavior is sufficiently caused by factors beyond their control (*i.e.*, genes and upbringing):

Imagine there is a world where the beliefs and values of every person are caused completely by the combination of one's genes and one's environment. For instance, one day in this world, two identical twins, named Fred and Barney, are born to a mother who puts them up for adoption. Fred is adopted by the Jerks and Barney is adopted by the Kinders. In Fred's case, his genes and his upbringing by the selfish Jerk family have caused him to value money above all else and to believe it is OK to acquire money however you can. In Barney's case, his (identical) genes and his upbringing by the kindly Kinder family have caused him to value honesty above all else and to believe one should always respect others' property. Both Fred and Barney are intelligent individuals who are capable of deliberating about what they do. One day Fred and Barney each happen to find a wallet containing \$1000 and the identification of the owner (neither man knows the owner). Each man is sure there is nobody else around. After deliberation, Fred Jerk, because of his beliefs and values, keeps the money. After deliberation, Barney Kinder, because of his beliefs and values, returns the wallet to its owner. Given that, in this world, one's genes and environment completely cause one's beliefs and values, it is true that if Fred had been adopted by the Kinders, he would have had the beliefs and values that would have caused him to return the wallet; and if Barney had been adopted by the Jerks, he would have had the beliefs and values that would have caused him to keep the wallet.

Despite what we took to be a robust description of complete causation by genes and environment, a significant majority of students (76 percent) judged both that Fred kept the wallet of his own free will and Barney returned it of his own free will. This response pattern was very similar to the pattern of students' judgments about free will in the Jeremy cases, suggesting that this scenario probed similar intuitions about the relationship between determinism and free will. We also tested whether students judge that Fred is "morally blameworthy for keeping the wallet" and that Barney is "morally praiseworthy for returning the wallet." For almost all students, these judgments were consistent, and the response patterns were not significantly different from the judgments we collected about free will: 60 percent judged that Fred is blameworthy and 64 percent judged that Barney is praiseworthy.

Finally, we ran a third survey in which we described determinism in terms of a universe being re-created over and over such that given the same initial conditions and laws of nature, everything must happen the same way every time, including an agent's stealing a necklace at a particular time every time the universe is re-created. Most students (66 percent) judged that the agent acts of her own free will, and most (77 percent) judged her to be morally responsible for her action. These results are consistent with those in the other two cases. They are *inconsistent*, however, with the incompatibilist prediction that most people recognize that "there is some kind of conflict between freedom and determinism" (Kane 1999).

Indeed, since incompatibilism is usually touted as the intuitive and commonsensical view, it is surprising that so many students judged that agents in a deterministic universe acted freely and are morally responsible for their actions. Not surprisingly, other experimental philosophers have developed studies to test whether there are alternative explanations to our data—for instance, whether there are conditions under which people are more likely to express incompatibilist intuitions (Nichols and Knobe 2007). This preliminary experimental work on free will and moral responsibility has generated an interesting discussion about how best to understand people's pre-philosophical intuitions and judgments and also about how empirically informed data about such intuitions and judgments should be—or should *not* be—used in philosophical debates (Nadelhoffer and Nahmias 2007; Kauppinen 2007; Knobe and Doris forthcoming; Sosa 2007).

As we've mentioned, experimental philosophers are also polling students (and other folk) to better understand ordinary intuitions and conceptual usage regarding, among other topics, knowledge, intentional action, moral judgments, and reference. We will not review that work here, but in the Appendix we offer some examples of the thought experiments used in some of these studies, along with references to some famous thought experiments in the literature, as a resource for those who are interested in using such surveys as a tool to teach philosophy. We will now explain in more detail *why* we think philosophy teachers will find polling their students in a systematic way to be a valuable tool and *how* we suggest teachers use it in the classroom.

Why Use Surveys in the Philosophy Classroom

We have found polling students to be an excellent example of good research going hand-in-hand with good teaching. During the course of running our studies on free will and determinism, we found that handing out surveys to students that probed their intuitions about thought experiments and philosophical problems not only generated some

interesting results, but it also worked remarkably well to get students to engage with the material. Whether we were lecturing about the Trolley Problem, Gettier cases, Frankfurt cases, Nozick's experience machine, and the like (see Appendix), polling has proven to be a useful pedagogical tool. So we started handing out surveys to students even when the results would not be included in any formal studies. A tool that we originally introduced into our classrooms for research purposes turned out to be a valuable pedagogical tool as well.⁹

Of course, using thought experiments in philosophy classes is nothing new—most of us do it all the time. So, why should we spend the extra time it takes to type up, copy, and hand out surveys to students when we can simply describe the thought experiment and poll them as a group? There are several benefits of using formal surveys as a way of introducing students to philosophical topics and issues.

1) *Clear and fair presentation of the thought experiment.* Many of the surveys you have students take will be based on thought experiments they will encounter in the readings. By surveying students, not only do you familiarize them with some of the puzzles that interest philosophers, but you also help foster student engagement with the material—which we have found is perhaps the main pedagogical benefit of polling. Rather than talking through the thought experiments and hoping both that you present them accurately and that the students keep up with them and remember their important details, they can have a written version in front of them to read carefully and refer back to as needed. Students can reflect on the cases sufficiently to give thoughtful responses to the questions that follow them. It can be useful to use the more rhetorical intuition pumps in the literature, since you will often assign the articles from which they derive to the students. And you can also discuss the thought experiment as a philosophical tool and ways to respond to this tool, including analysis of the subtly misleading—or at least leading—elements of the intuition pump. But it can also be useful to use the scenarios developed by experimental philosophers (see Appendix), since these are generally designed both to be clear and understandable by laypersons and also to offer a fair presentation of the issues. Of course, you can also develop your own version of the thought experiments! (See note 2 for references to collections of thought experiments.)

2) *Clear and accurate presentation of the polling results.* When philosophy teachers poll their students, they usually have them raise their hands and take a quick tally of the results, “guesstimating” in terms of rough proportions—“OK, it looks like we have one quarter who are consistent utilitarians and three quarters who are utilitarians unless they have to push the fat man to his death to stop the runaway trolley.” It is more useful to tabulate their responses more accurately and present

them clearly on the board for students to see. They will often be quite surprised to see how many students do not share their own intuitions, and this will help spark discussion, with the students realizing they need to come up with some justifications for their own beliefs regarding the cases since their beliefs are not shared by many other students.

3) *Suggests to students that their opinions are important.* Having students participate in a class survey is a nice way of showing them that their opinions are interesting and important enough for philosophers to pay attention to—even if the results of the poll are not going to end up being part of any formal studies. As a result, we have found that students end up being more likely to carefully read the vignettes and respond to the questions. By giving them five minutes at the beginning of class to quietly reflect on a philosophical puzzle, you give students the chance to do what we philosophers do all the time. Indeed, we have been surprised by how many students took the time to carefully explain their answers to the survey questions.

4) *Initiates discussion and debate.* If you simply ask students out loud what they think about a thought experiment, they have far less time to think about the problem. As a result, fewer students will be comfortable explaining their answers in front of their classmates. If you give them a few minutes to quietly puzzle over a philosophical problem, on the other hand, you make it more likely that they will commit to a particular answer or position and feel more comfortable discussing their views since they have had some time to flesh out their ideas and intuitions. Students who feel more strongly about a philosophical problem are more likely to participate in class discussions about that problem—especially once students are informally divided into “yes” and “no” groups (you can also divide them more formally, putting them on opposite sides of the classroom to debate the topic). Moreover, students with similar intuitions are more likely to support one another’s arguments and suggestions, thereby further increasing the willingness of others to participate. In our experience, surveying students increases how invested they end up feeling about a philosophical topic or problem. Finally, we always have students write out on the back of the surveys their explanations for why they made the judgments they did. This leads them to begin to try to find justifications for their philosophical intuitions. The subtle difference between writing down one’s opinion and why one holds it, rather than just considering one’s answer and raising one’s hand, can also change how invested the students feel about their views.

5) *Allows you to get to know your students and allows them to get to know each other.* Another benefit of using formal surveys in the classroom is that doing so affords the teacher an opportunity to get a more accurate feel for how a particular group of students responds

to particular thought experiments. Knowing in advance where your students are in terms of their attitudes and judgments concerning a philosophical issue enables you to focus more on things you know many of them find to be counter-intuitive. Not only does surveying students provide teachers with more insight into their students' beliefs and judgments, it also gives teachers the chance to see whether their students' intuitions about a particular issue change during the course of the class.¹⁰ For instance, if students initially have consequentialist or naïve relativist intuitions, how likely is it that they will be persuaded by Kant's *arguments*? What about thought experiments that purportedly *show* rather than *say* what's wrong with consequentialism or naïve relativism? Do students find these intuition pumps to be more persuasive than the lengthy philosophical arguments that contain them? If so, that would be a curious and important thing to know. This represents yet another way in which surveying students produces positive pedagogical side effects.

Formally surveying students and having them express their responses also allows them to get to know what their peers think about the issues. Finally, teachers can also let their students know what they think. Though there are important issues about neutrality vs. advocacy to consider (see, *e.g.*, Bomstad 1995), in some cases teachers may decide to explain why they answer the survey questions in a particular way. And they may explain why they initially had certain intuitions about cases but, through reflective equilibrium or theorizing, came to view those intuitions as mistaken.

How to Use Surveys in the Classroom

- Pass the surveys out *before* any discussion or reading on the topic (it is often best to hand them out at the beginning of class¹¹ but sometimes good to have them do them at end of class or as homework to think about before next class, or before they do the relevant reading).
- With larger classes you can have students help you pass the surveys out, take them up, and quickly tally the results.
- We usually write at the top of the survey something like, "Please read the scenario below carefully. Then answer the questions that follow. There are no right or wrong answers. I am interested in finding out what your intuitions about the scenario are."
- It is crucial to use scenarios that are clear and understandable to the students. This means that you sometimes need to modify ones that are used in the philosophical literature. Try to avoid any technical language or philosophical terms of art. You may sometimes want to use slightly different versions of scenarios with different students to see what effects

the differences have (*e.g.*, the bystander at the switch and footbridge versions of the trolley problem). Occasionally it can be useful to have students answer questions to a series of related scenarios.

- Just after the scenario, if appropriate, you may want to include a reminder to think conditionally or counterfactually (*e.g.*, “please answer the following questions on the assumption that the scenario is true,” “even if you think this scenario is implausible, imagine it is the case as you answer the following questions”). This also opens an opportunity to discuss conditional and counterfactual reasoning with the students (see note 8).
- For the questions, it’s usually best to use “yes,” “no,” and “I don’t know” as the possible answers but you can also use a Likert scale (*e.g.*, “to what degree do you agree or disagree” with a -3 to +3 scale). It’s usually best not to use too many questions. It can be interesting to have different students answer different questions about the same scenario to find correlations and conflicts.
- Finally, on the back of the survey (and only *after* they have answered all the questions) you should have students write a short explanation of why they answered the questions as they did. As discussed above, this will get them thinking more deeply about their reasons for having the intuitions they have and make them more comfortable expressing their opinion if you ask for volunteers or call on them. And it will also allow you to get a better sense of your students’ thoughts on the issues you are teaching.
- Again, put the results on the board. They will almost never be unanimous. If they are closely split, you can have the students debate. If the results are lop-sided, the majority is usually very interested to hear how the minority could possibly think the crazy things they do. And members of the minority are often happy to oblige.

Conclusion

Experimental philosophy represents an interesting—and we believe productive—new methodology in academic philosophy. Even if it turns out to be less important for research than many experimental philosophers have assumed, the formal survey methodology it employs will remain a valuable tool for teaching philosophy for the reasons we have discussed in this paper. We firmly believe that surveying students is one of the easiest and most effective ways to get one’s students to start actively philosophizing rather than merely passively learning philosophy. Minimally, by giving students surveys that afford them with the opportunity to quietly flesh out the details of philosophical thought experiments and then having them discuss these experiments with one another in class, we make it more likely that they will engage in the kind of Socratic dialogue that is the lifeblood of philosophy.

Appendix: Some Examples of Scenarios for Surveying Students

- The Knobe Effect (Knobe 2003) regarding Intentional Action (different subjects get different scenarios with one of the two words in the brackets):

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

Question: Did the chairman intentionally [harm, help] the environment?

(Knobe and others have found that a robust majority responds that he did intentionally *harm* the environment while a robust majority responds that he did *not* intentionally *help* the environment.)

- The Butler Problem (1977)—Similar scenario used in Nadelhoffer 2004:

If Brown in an ordinary game of dice hopes to throw a six and does so, we do not say that he threw the six intentionally. On the other hand if Brown puts one cartridge into a six-chambered revolver, spins the chamber as he aims it at Smith and pulls the trigger hoping to kill Smith, we would say if he succeeded that he had killed Smith intentionally. How can this be so, since the probability of the desired result is the same?

- Gettier Cases (Gettier 1963)—Scenarios used in Nichols, Stich, and Weinberg 2003:

Case 1: Bob has a friend, Jill, who has driven a Buick for many years. Bob therefore thinks that Jill drives an American car. He is not aware, however, that her Buick has recently been stolen, and he is also not aware that Jill has replaced it with a Pontiac, which is a different kind of American car.

Question: Does Bob really know that Jill drives an American car, or does he only believe it?

Case 2: Pat is at the zoo with his son, and when they come to the zebra cage, Pat points to the animal and says, “that’s a zebra.” Pat is right—it is a zebra. However, given the distance the spectators are from the cage, Pat would not be able to tell the difference between a real zebra and a mule that is cleverly disguised to look like a zebra. And if the animal had really been a cleverly disguised mule, Pat still would have thought that it was a zebra.

Question: Does Pat really know that the animal is a zebra, or does he only believe that it is?

- Brain-in-a-Vat Scenario—Scenario used in Nichols *et al.* 2003:

Scenario: George and Omar are roommates, and enjoy having late-night ‘philosophical’ discussions. One such night Omar argues, “At some point in time, by, like, the year 2300, the medical and computer sciences will be able to simulate the real world very convincingly. They will be able to grow a brain without a body, and hook it up to a supercomputer in just the right way so that the brain has experiences exactly as if it were a real person walking around in a real world, talking to other people, and so on. And so the brain would believe it was a real person walking around in a real world, *etc.*, except that it would be wrong—it’s just stuck in a virtual world, with no actual legs to walk and with no other actual people to talk to. And here’s the thing: how could you ever tell that it isn’t really the year 2300 now, and that you’re not really a virtual-reality brain? If you were a virtual-reality brain, after all, everything would look and feel exactly the same to you as it does now!” George thinks for a minute, and then replies: “But, look, here are my legs.” He points down to his legs. “If I were a virtual-reality brain, I wouldn’t have any legs really—I’d only really be just a disembodied brain. But I know I have legs—just look at them!—so I must be a real person, and not a virtual-reality brain, because only real people have real legs. So I’ll continue to believe that I’m not a virtual-reality brain.” George and Omar are actually real humans in the actual real world today, and so neither of them are virtual-reality brains, which means that George’s belief is true.

Question: Does George know that he is not a virtual-reality brain, or does he only believe it?

- Robert Nozick’s Experience Machine (Nozick 1974):

Suppose there were an experience machine that would give you any experience you desired. Super-duper neuropsychologists could stimulate your brain so that you would think and feel you were writing a great novel, or making a friend, or reading an interesting book. All the time you would be floating in a tank, with electrodes attached to your brain. Should you plug into this machine for life, preprogramming your life’s desires? Of course, while in the tank you won’t know that you’re there; you’ll think it’s all actually happening. Others can also plug in to have the experiences they want, so there’s no need to stay unplugged to serve them (ignore problems such as who will service the machines if everyone plugs in).

Question: Would you plug in?

- Bernard Williams’s Jim and the Indians (Williams 1973):

Jim finds himself in the central square of a small South American town. Tied up against the wall are a row of twenty Indians, most ter-

rified, a few defiant, in front of them several armed men in uniform. A heavy man . . . turns out to be the captain in charge and, after a good deal of questioning of Jim which established that he got there by accident while on a botanical expedition, explains that the Indians are a random group of the inhabitants who, after recent acts of protest against the government, are just about to be killed to remind other possible protestors of the advantages of not protesting. However, since Jim is an honored visitor from another land, the captain is happy to offer him a guest's privilege of killing one of the Indians himself. If Jim accepts, then as a special mark of the occasion, the other Indians will be let off. Of course, if Jim refuses, then there is no special occasion, and Pedro here will do what he was about to do when Jim arrived, and kill them all.

Question: Should Jim kill one Indian to spare the lives of nineteen others?

- Judith Jarvis Thomson's Trolley Problem (1985)—Similar scenario used in Greene *et al.* 2001:

Suppose you are the driver of a trolley. The trolley rounds a bend, and there comes into view ahead five track workmen, who have been repairing the track. The track goes through a bit of a valley at that point, and the sides are steep, so you must stop the trolley if you are to avoid running the five men down. You step on the brakes, but, alas, they don't work. Now you suddenly see a spur of track leading off to the right. You can turn the trolley onto it, and thus save the five men on the straight track ahead. Unfortunately, Mrs. Foot has arranged that there is one track workman on that spur of track. He can no more get off the track in time than the five can, so you will kill him if you turn the trolley onto the spur.

Question: Is it morally permissible for you to turn the trolley?

- Thomson's Society of Music Lovers (Thomson 1971):

Imagine this. You wake up and find yourself back to back in bed with an unconscious violinist. A famous unconscious violinist. He has been found to have a fatal kidney ailment, and the Society of Music Lovers has canvassed all the available medical records and found that you alone have the right blood type to help. They have therefore kidnapped you, and last night the violinist's circulatory system was plugged into yours, so that your kidneys can be used to extract poisons from his blood as well as your own. The director of the hospital now tells you, "Look, we're sorry the Society of Music Lovers did this to you . . . but still, they did it, and the violinist is now plugged into you. To unplug you would be to kill him. But never mind, it's only for nine months. By then he will have recovered from his ailment, and can be safely unplugged from you.

Question: Is it morally permissible for you to simply unplug yourself from the violinist right away?

- Locke's Locked Room (Locke 1991):

Suppose a man is carried, while fast asleep, into a room, where there is a person he longs to see and speak with: and suppose he is locked in the room, beyond his power to get out: he awakes, and is glad to find himself in so desirable company, which he stays willingly in, *i.e.*, prefers his stay to going away.

Questions: Does the man stay in the room (a) voluntarily, and (b) freely?

- Frankfurt's Nefarious Surgeon (Frankfurt 1969)—Related scenarios used in Woolfolk, Doris, and Darley 2006:

Suppose someone—Black, say—wants Jones to perform a certain action. Black is prepared to go to considerable length to get his way, but he prefers to avoid showing his hand unnecessarily. So he waits until it is clear to him (Black is an excellent judge of such things) that Jones is going to decide to do something other than what he wants him to do. Whatever Jones's initial preferences and inclinations then, Black will have his way. . . . Now suppose Black never had to show his hand because Jones, for reasons of his own, decides to perform and does perform the very action Black wants him to perform.

Questions: Is Jones responsible for performing the action in question even though he could not have avoided performing it? Does Jones freely perform the action even though he could not have avoided performing it?

- Lehrer's Truetemp Case (Lehrer 1990)—Scenario used in Swain, Alexander, and Weinberg 2006:

Suppose a person, whom we shall name Mr. Truetemp, undergoes brain surgery by an experimental surgeon who invents a small device which is both a very accurate thermometer and a computational device capable of generating thoughts. The device, call it a tempucomp, is implanted in Truetemp's head so that the very tip of the device, no larger than the head of a pin, sits unnoticed on his scalp and acts as a sensor to transmit information about the temperature to the computational system of his brain. This device, in turn, sends a message to his brain causing him to think of the temperature recorded by the external sensor. Assume that the tempucomp is very reliable, and so his thoughts are correct temperature thoughts. All told, this is a reliable belief-forming process. Now imagine, finally, that he has no idea that the tempucomp has been inserted in his brain, is only slightly puzzled about why he thinks so obsessively about the temperature, but never checks a thermometer to determine whether these thoughts about the

temperature are correct. He accepts them unreflectively, another effect of the tempucomp. Thus, he thinks and accepts that the temperature is 104 degrees. It is.

Question: Does he know that it is?

- Nichols's Psychopathic Killer (Nichols 2004c):

John is a psychopathic criminal. He is an adult of normal intelligence, but he has no emotional reaction to hurting other people. John has hurt and indeed killed other people when he has wanted to steal their money. He says that he knows that hurting others is wrong, but that he just doesn't care if he does things that are wrong.

Question: Does John really understand that hurting others is morally wrong?

- Kripke (1972) cases of Reference—Scenario used in Machery *et al.* (2004):

Suppose that John has learned in college that Gödel is the man who proved an important mathematical theorem, called the incompleteness of arithmetic. John is quite good at mathematics and he can give an accurate statement of the incompleteness theorem, which he attributes to Gödel as the discoverer. But this is the only thing that he has heard about Gödel. Now suppose that Gödel was not the author of this theorem. A man called "Schmidt," whose body was found in Vienna under mysterious circumstances many years ago, actually did the work in question. His friend Gödel somehow got hold of the manuscript and claimed credit for the work, which was thereafter attributed to Gödel. Thus, he has been known as the man who proved the incompleteness of arithmetic. Most people who have heard the name "Gödel" are like John; the claim that Gödel discovered the incompleteness theorem is the only thing they have ever heard about Gödel.

Question: When John uses the name "Gödel," is he talking about:

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

(B) the person who got hold of the manuscript and claimed credit for the work?

Notes

We would like to thank the experimental philosophers we have cited (and others we have not) for their interesting work and the surveys they have developed that we can use to teach our students. We especially thank Jason Turner and Steve Morris, our co-authors on our papers surveying folk intuitions about free will and moral responsibility.

1. It could be that people don't ordinarily have any beliefs or intuitions about some philosophical issues, such as mereology or set theory, but they are likely to have some

salient beliefs and intuitions about issues such as human knowledge, morality, intentional action, free will, moral responsibility, and legal culpability.

2. Theodore Schick and Lewis Vaughn (2002) have put together an interesting textbook that relies heavily on thought experiments as an effective way of teaching philosophy. The textbook—when combined with the formal surveying methods we discuss in this paper—would likely make for an effective way of getting students to learn philosophy. See also Tittle (2005) and Baggini (2005) for extensive collections of thought experiments.

3. Perhaps the best example of how conceptual analysis so conceived is supposed to work is the age-old definition of knowledge as justified true belief. According to Edmund Gettier, at least, this definition simply won't do (Gettier 1963: 121–22), since he provides cases involving agents who seemingly have justified true beliefs that *x*, but who, *intuitively*, do *not* know that *x*.

4. Determining precisely what intuitions are is a tricky matter that would take us too far afield for present purposes. For now, we are simply going to follow Alvin Goldman and Joel Pust in suggesting that, “the contents of intuitions are usually singular classificational propositions, to the effect that such-and-such an example is or is not an instance of knowledge, of justice, of personal identity, and so forth” (Goldman and Pust 1998: 182). So, students' reflective responses to questions about well-constructed scenarios describing philosophical issues are generally a fair indication of their intuitions about those issues.

5. One natural response for philosophers whose positions do not settle with folk intuitions would be to explain why these intuitions are mistaken (*i.e.*, offer an error theory) or why these folk concepts need to be revised. Alternatively, they might offer an explanation for why we have the intuitions we do but why they do not in fact commit us to certain conceptual or theoretical views. But these moves still require that we first determine—rather than merely speculate about—what these folk intuitions and concepts actually are. Otherwise, we will not know exactly what it is that needs to be explained away or revised. This suggests that we must first make an earnest attempt to probe and describe the folk intuitions in question before subsequent philosophical analyses are developed.

6. For an in depth examination of some of the meta-philosophical and methodological issues that arise in experimental philosophy see Kauppinen (2007) and Nadelhoffer and Nahmias (2007).

7. For details of the following studies, including more on methodology, statistics, and philosophical implications, see Nahmias *et al.* 2005 and 2006.

8. In pilot studies we found that some participants seemed to fail to reason conditionally (*e.g.*, given their explanations on the back of the survey, some seemed to assume that the scenario is impossible *because* Jeremy has free will, rather than making judgments about Jeremy's freedom on the assumption that the scenario is actual). To correct for this problem, we used manipulation checks to determine comprehension and excluded participants who missed the checks, and we used an initial question which asked participants whether they think the scenario is possible (the majority responded ‘no,’ offering various reasons on the back of the survey). Then they were then asked to ‘suspend disbelief’ for the experimental question concerning whether Jeremy acted of his own free will, *etc.* Participants were specifically told that regardless of how they answered the question 1, they were to *imagine* that such a supercomputer actually did exist and actually could predict the future, including Jeremy's robbing the bank (and they were told to assume that Jeremy does not know about the prediction). Whether people can really suspend belief in this way is both an interesting and open question. It also raises interesting issues about how to teach students to reason conditionally and counterfactually.

9. Though experimental philosophers who plan to publish their data should get approval from their institution's Human Subjects Committee (or IRB) to carry out such

research, using written surveys in the classroom for pedagogical purposes does not require such approval and does not appear to raise any special ethical concerns.

10. Though you may want your students' responses to remain anonymous (as experimental philosophers do in their formal surveys), you may sometimes want students to turn in their answers with their names. This seems unproblematic, since we often ask students to present and defend their opinions about philosophical issues.

11. We suggest handing out surveys in class, but another method that might be worth experimenting with is to post surveys online at your course website and have students take them for homework. There are several online survey programs (such as www.questionpro.com and www.surveymonkey.com) that offer features such as tabulation and statistical analyses of results.

Bibliography

- Adams, F., and A. Steadman. 2004a. "Intentional Action in Ordinary Language: Core Concept or Pragmatic Understanding?" *Analysis* 64: 173–81.
- . 2004b. "Intentional Actions and Moral Considerations: Still Pragmatic," *Analysis* 64: 264–68.
- Baggini, J. 2005. *The Pig that Wants to be Eaten: 100 Experiments for the Armchair Philosopher*. New York: Penguin.
- Bomstad, L. 1995. "Advocating Procedural Neutrality," *Teaching Philosophy* 18(3): 197–210.
- Brant, R. B. 1959. *Ethical Theory*. Englewood Cliffs, N.J.: Prentice Hall.
- Butler, R. 1978. Report on *Analysis* "problem" no. 6. *Analysis* 38: 113–14.
- DePaul, M. R., and W. Ramsey, eds. 1998. *Rethinking Intuition: The Psychology of Intuition And its Role in Philosophical Inquiry*. Lanham, Md.: Rowman and Littlefield.
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23: 121–22.
- Goldman, A., and J. Pust. 1998. "Philosophical Theory and Intuitional Evidence," in DePaul and Ramsey 1998: 179–97.
- Frankfurt, H. 1969. "Alternative Possibilities and Moral Responsibility," *Journal of Philosophy* 66: 828–39.
- Greene, J. D., R. B. Sommerville, L. E. Nystrom, J. M. Darley, and J. D. Cohen. 2001. "An fMRI Investigation of Emotional Engagement in Moral Judgment," *Science* 293: 2105–08.
- Hoff-Sommer, C., ed. 1986. *Right and Wrong*. New York: Harcourt, Brace, Jovanovich.
- Kane, R. 1999. "Responsibility, Luck, and Chance: Reflections on Free Will and Indeterminism," *Journal of Philosophy* 96: 217–40.
- Kauppinen, A. 2007. "The Rise and Fall of Experimental Philosophy," *Philosophical Explorations* 10(2): 95–118.
- Knobe, J. 2003. "Intentional Action and Side Effects in Ordinary Language," *Analysis* 63: 190–94.
- . 2004. "Intention, Intentional Action, and Moral Considerations," *Analysis* 64: 181–87.
- Knobe, J., and J. Doris. Forthcoming. "Strawsonian Variations: Folk Morality and the Search for a Unified Theory," in *The Handbook of Moral Psychology*, ed. J. Doris, G. Harman, S. Nichols, J. Prinz, W. Sinnott-Armstrong, and S. Stich. Oxford: Oxford University Press.
- Kripke, S. 1972/1980. *Naming and Necessity*. Cambridge, Mass.: Harvard University Press.
- Lehrer, K. 1990. *Theory of Knowledge*. Boulder, Colo.: Westview Press.

- Locke, J. 1991. *An Essay Concerning Human Understanding*. Oxford: Clarendon Press.
- Machery, E., R. Mallon, S. Nichols, and S. Stich. 2004. "Semantics, Cross-Cultural Style," *Cognition* 92: B1–B12.
- Malle, B. F., L. J. Moses, and D. A. Baldwin, eds. 2001. *Intentions and Intentionality*. Cambridge, Mass.: MIT Press.
- McCann, Hugh. 1998. *The Works of Agency: On Human Action, Will, and Freedom*. Ithaca, N.Y.: Cornell University Press.
- Mele, A. 1992. "Intentions, Reasons, and Beliefs: Morals of the Toxin Puzzle," *Philosophical Studies* 68: 171–94.
- _____. , ed. 1997. *The Philosophy of Action*. Oxford: Oxford University Press.
- _____. . 2001. "Acting Intentionally: Probing Folk Notions," in Malle, Moses, and Baldwin 2001.
- Nadelhoffer, T. 2004. "The Butler Problem Revisited," *Analysis* 64(3): 277–84.
- _____. . 2005. "Skill, Luck, Control, and Intentional Action," *Philosophical Psychology* 18(3): 343–54.
- _____. . 2006. "Bad Acts, Blameworthy Agents, and Intentional Actions: Some Problems for Jury Impartiality," *Philosophical Explorations* 9(2): 203–20.
- Nadelhoffer, T., and E. Nahmias. 2007. "The Past and Future of Experimental Philosophy," *Philosophical Explorations* 10(2): 123–49.
- Nahmias, E., S. Morris, T. Nadelhoffer, and J. Turner. 2005. "Surveying Free Will: Folk Intuitions about Free Will and Moral Responsibility," *Philosophical Psychology* 18(5): 561–84.
- _____. . 2006. "Is Incompatibilism Intuitive?" *Philosophy and Phenomenological Research* 73(1): 28–53.
- Nichols, S. 2004a. "After Objectivity," *Philosophical Psychology* 17: 5–28.
- _____. . 2004b. "The Folk Psychology of Free Will: Fits and Starts," *Mind & Language* 19(5): 473–502.
- _____. . 2004c. *Sentimental Rules*. Oxford: Oxford University Press.
- Nichols, S., and J. Knobe. 2007. "Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions," *Nous* 41: 663–85.
- Nichols, S., S. Stich, and J. Weinberg. 2003. "Metaskepticism: Meditations in Ethno-Epistemology," in *The Skeptics*, ed. S. Luper. Burlington, Vt.: Ashgate, 227–47.
- Nozick, R. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.
- O'Connor, T. 2000. *Persons and Causes: The Metaphysics of Free Will*. Oxford: Oxford University Press.
- Pereboom, D. 2001. *Living Without Free Will*. Cambridge: Cambridge University Press.
- Ramsey, W. 1998. "Prototypes and Conceptual Analysis," in DePaul and Ramsey 1998: 161–77.
- Schick, T., and L. Vaughn. 2002. *Doing Philosophy: An Introduction through Thought Experiments*. New York: McGraw-Hill.
- Sosa, Ernest. 2007. "Experimental Philosophy and Philosophical Intuition," *Philosophical Studies* 132(1): 99–107.
- Stich, S., and J. Weinberg. 2001. "Jackson's Empirical Assumptions," *Philosophy and Phenomenological Research* 52: 637–42.
- Swain, S., J. Alexander, and J. Weinberg. 2006. "The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp." *Online Philosophy Conference*.
- Thomson, J. J. 1971. "A Defense of Abortion," *Philosophy and Public Affairs* 1: 47–66.
- _____. . 1985. "The Trolley Problem," *Yale Law Journal* 94: 1395–1415.
- Tittle, P. 2005. *What If . . . : Collected Thought Experiments in Philosophy*. New York: Pearson.

Williams, B. 1973. "A Critique of Utilitarianism," in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against*. Cambridge: Cambridge University Press, 1973, 82–117.

Woolfolk, R., J. Doris, and J. Darley. 2006. "Identification, Situational Constraint, and Social Cognition: Studies in the Attribution of Moral Responsibility," *Cognition* 100: 283–301.

Thomas Nadelhoffer, Philosophy, Dickinson College, Carlisle, PA tadelhoffer@gmail.com; Eddy Nahmias, Philosophy, Georgia State University, Atlanta, GA 30302; phlean@langate.gsu.edu