

How to Do Better: Toward Normalizing Experimentation in Epistemology*

John Turri

john.turri@gmail.com

Abstract: Appeals to ordinary thought and talk are frequent in philosophy, perhaps nowhere more than in contemporary epistemology. When an epistemological theory implies serious error in “commonsense” or “folk” epistemology, it is counted as a cost of the view. Similarly, when an epistemological theory respects or vindicates deep patterns in commonsense epistemology, it is viewed as a benefit of the view. Philosophers typically rely on introspection and anecdotal social observation to support their characterizations of commonsense epistemology. But recent experimental research shows that philosophers employing these methods often seriously mischaracterize commonsense. Based on these findings, I propose a fundamental change to standard practice in the field. Whether the fundamental goal of epistemology is descriptive or prescriptive, experimentation is an integral part of the project.

Keywords: philosophical method; experimental epistemology; folk epistemology; knowledge attributions; knowledge; evidence; belief; truth

Word count: 4731

* This is the penultimate version of a paper to appear in J. Nado (Ed.), *Advances in experimental philosophy and philosophical methodology* (Bloomsbury Academic).

Introduction

Appeals to ordinary thought and talk are frequent in philosophy. For example, Aristotle claimed that a theory of happiness should be evaluated in light of “what is commonly said about it” (Aristotle 350BCE, 1098b, 9-11). He judged it improbable that a widely held view would be completely incorrect (1098b, 28-29). John Locke claimed that a theory of knowledge should be informed by how we ordinarily talk about knowledge (Locke 1690, book 4.11.3-8). Thomas Reid claimed, “Philosophy has no other root but the principles of Common Sense,” and that in a contest between “Common Sense and Philosophy, the latter will always come off both with dishonour and loss” (Reid 1764, p. 19). More recently, J.L. Austin advised that “ordinary language” should get “the first word” in philosophical theorizing (Austin 1956, p. 11). And Wilfrid Sellars argued that identifying the defining features of ordinary thought — “the manifest image” — is “a task of the first importance” for philosophers (Sellars 1963, ch. 1). Other leading philosophers, including G.E. Moore and Roderick Chisholm, held common sense and ordinary language in similarly high regard (Lemos 2004).

Many lines of research in contemporary epistemology are based on assumptions about ordinary thought and talk about knowledge — that is, about alleged central tendencies in commonsense epistemology. In the next section, I review five case studies where epistemologists have seriously mischaracterized commonsense epistemology. In the Conclusion, I consider some implications of these findings. Epistemology, as currently practiced, fails in its ambitions and might be alienating people from the field. Consequently, standard practice in the field should change.

The good news is that the tools for change are already at the field's disposal.

Case Studies

Knowledge and reliability

Contemporary philosophers nearly unanimously agree that knowledge requires reliability (for references, see Turri in press a). On this view, knowledge must be produced by abilities that “will or would yield mostly true beliefs” (Alston 1995: 8-9). An enormous amount of work has been based on the assumption that knowledge requires reliability (for recent reviews, see Goldman 2008; Becker 2009). But there are theoretical reasons to expect that knowledge does not require reliability (Turri in press a; Turri in press b), and the literature contains surprisingly little explicit argumentation for reliabilism about knowledge (hereafter just “reliabilism”). Instead, reliabilists typically appeal to commonsense epistemology to support reliabilism. They claim that reliabilism best fits the patterns in “ordinary, intuitive judgments” about knowledge (Dretske 1981, p. 92; see also p. 249, n. 8), that reliabilism is just “commonsense” (Sosa 1991, p. 100), and that our “inclinations” and “intuitions” about knowledge are keyed to reliability (Goldman 1993, p. 271).

This “proto-reliabilist hypothesis,” as we might call it, has testable predictions. First and foremost, people will not count unreliably produced belief as knowledge. For example, if an ability produces only ten percent or thirty percent true beliefs, then the beliefs it produces will not be

judged knowledge. Second, and closely related, vast and explicit differences in reliability will produce large differences in knowledge judgments.

It is easy to test these predictions: we just need to elicit judgments about simple cases that vary how reliably beliefs are formed. When this was done, the results were grim for the proto-reliabilist hypothesis (Turri in press e). Commonsense fully embraces the possibility of unreliable knowledge: rates of unreliable knowledge attribution sometimes reach 90%. Knowledge judgments are surprisingly insensitive to information about reliability: people attribute knowledge at similar rates whether the agent gets it right ninety percent of the time or ten percent of the time.

For example, in one study, participants were divided into two groups. They all read a story about Alvin, who has been out running errands all day. In the morning, Alvin's wife told him that he should stop at the dry cleaners. On his way home, Alvin stops at the dry cleaners. Participants in the unreliable condition were also told that Alvin's memory is very "unreliable"; participants in the reliable condition were told that Alvin's memory is very "reliable." That was the only difference. Participants were then asked whether Alvin knows that he should stop at the dry cleaners. In the unreliable condition, participants categorized Alvin as unreliable and 86% of them attributed knowledge; in the reliable condition, participants categorized Alvin as reliable and 88% of them attributed knowledge (Turri in press e). Similar results were observed for different cognitive faculties, ways of describing unreliability, and narrative contexts.

Knowledge and context

Epistemic contextualism is the view that the verb “know” is a context sensitive expression. More specifically, according to contextualism, in order for us to truthfully say a person “knows” a proposition, that person must meet the evidential standard set by our context and, critically, the standard changes across contexts. In support of their view, contextualists claim that it best explains competent speakers’ linguistic behavior in certain situations (Lewis 1996; DeRose 2009; Cohen 2013). That is, contextualism has its “basis in ordinary language,” where we allegedly find “evidence of the very best type” that the standards of “knowledge” are context sensitive (DeRose 2009, p. 48). Contextualism is motivated on the grounds that it is uniquely positioned to explain “evident facts about usual and . . . appropriate linguistic behavior” (DeRose 2009, p. 81).

The most famous way of illustrating the idea involves a pair of cases about a man who wants to deposit a check and is deciding whether to wait in a long line at the bank on a Friday afternoon, or come back on Saturday morning when the line would be short (DeRose 2009). But then the question arises: is this bank actually open tomorrow (Saturday) morning? The man visited this bank two Saturdays ago and it was open then, but banks do sometimes change their hours. In the “low stakes” version of the case, nothing serious hinges on whether he deposits the check before the weekend is over, and the man says, “I know that the bank is open tomorrow.” In the “high stakes” version of the case, something very serious hinges on whether he deposits the check before the weekend is over, and the man says, “I don’t know that the bank is open tomorrow.” Contextualists claim competent speakers will judge that the man truthfully says he

“knows” in the low stakes version, and that the man truthfully says he “doesn’t know” in the high stakes version.

It is an empirical question whether people behave as contextualists predict. Prior work has explored this empirical question, with mixed results. Some results seemed inconsistent with contextualist predictions (e.g. Buckwalter 2010; Feltz & Zarpentine 2010; Buckwalter 2014; see also May, Sinnott-Armstrong, Hull & Zimmerman 2010), but others were consistent with the predictions (e.g. Hansen & Chemla 2013; see also Alexander, Gonnerman & Waterman 2014). Contextualists have made several methodological objections to previous studies, particularly emphasizing the importance of eliciting metalinguistic judgments about a knowledge statement’s truth value and of having the agent say “I know” in the low case and “I don’t know” in the high case (DeRose 2009, p. 49, n. 2; DeRose 2011).

It turns out that, when the cases are tested while respecting those methodological strictures, the basic contextualist prediction turns out true: people tend to agree that the agent who says “I know” in the low stakes case speaks truthfully, and they tend to agree that the agent who says “I don’t know” in the high stakes case speaks truthfully (Turri in press f). Is this result good evidence that “know” is a context-sensitive expression?

No, it is not good evidence, because simpler alternative explanations immediately suggest themselves. On the one hand, across the pair of cases, people might draw different inferences about whether the agent satisfies requirements of knowledge that contextualists and non-contextualists alike accept, such as truth, belief, and evidence. For example, people might assume that someone who says “I don’t know” is less likely to have the relevant belief, or to have a *true* be-

lief, or to have good evidence than someone who says “I do know.” This could lead people to judge that the agent knows in the one case but not in the other, which in turn would lead people to judge that the one agent truthfully says “I know,” whereas the other agent truthfully says “I don’t know.” Contextualists do not accept these alternative explanations. They claim that it is “natural” to interpret the agent as equally confident in the low and high stakes cases (DeRose 2009, pp. 190-3). The claim that “nothing changes” about the quality of the agent’s evidence across the cases (Lawlor 2013, p. 81; see also DeRose 2009, p. 2). And some imply that we can or will “assume” that the relevant proposition is true in each case (DeRose 2009, p. 2).

Such interpretations might seem natural to contextualists, but it is an empirical question whether the cases are actually interpreted that way. And it is not hard to test the matter: we just need to (i) collect judgments about belief, truth, and evidence in the same context as the metalinguistic “knowledge” judgments, and (ii) see which judgments differ between low and high conditions. When this was tested, people judged that the low stakes agent is less likely to have the relevant belief, less likely to have a *true* belief, and less likely to have good evidence than the high stakes agent (Turri in press f). Any of or all of these differences could produce the observed pattern of metalinguistic judgments about “knowledge,” *even if* contextualism about “knows” is false. Similar results were observed for another famous pair of low/high cases proposed by contextualists (Cohen 1999).

On the other hand, the low/high cases contain a critical confound. Contextualists note that it is important for the agent in the low stakes case to say “I know,” and for the agent in the high stakes case say “I don’t know.” As a result, the low/high difference is confounded with the “I

know/I don't know" difference. Thus, an alternative, non-contextualist explanation of the finding is that people tend to defer to others' self-regarding knowledge statements. That is, people might simply assume that others are well positioned to report on whether they know a proposition, which could produce the pattern. This assumption could lead them to *defer* to the agent's self-regarding knowledge statement. If this *deferral hypothesis* is true, then we would expect the observed pattern of behavior even if contextualism is false.

Testing the deferral hypothesis is not difficult. The basic idea is to divide people into two groups and have them read the exact same story, except for one small difference: have the agent say "I do know" in one story, and "I don't know" in the other. Do not vary the stakes or anything else. When people were tested in this way, people exhibited the exact same pattern noted above: they agreed when the person said "I do know," and they agreed when the person said "I don't know" (Turri in press f). Similar results were observed for another famous pair of low/high cases proposed by contextualists (Cohen 1999).

The critical point here is that the difference between "I do know" and "I don't know" suffices, all by itself, to produce the pattern that contextualists predict. In other words, deferral is enough to produce the principal datum used to motivate contextualism. Varying the stakes is superfluous, as is the hypothesis that "know" is context-sensitive.

Knowledge and closure

In its simplest form, the epistemic closure principle states that, necessarily, if you know that P, and you know that Q if P, then you also know that Q. Or, at least, you would then be in a position

to easily know Q by simple inference. For instance, if you know that your car is parked in the parking structure, and you know that your car has not been stolen if it is parked in the structure, then you know that your car has not been stolen. The epistemic closure principle is implicated in many debates in contemporary epistemology, including debates about skepticism (Stroud 1984), the semantics of knowledge attributions (DeRose 1995; Cohen 1999), and the relationship between knowledge and practical considerations (Hawthorne 2004).

The epistemic closure principle is not as widely or uncritically accepted as, say, knowledge reliabilism, but it is very widely accepted. And, as with knowledge reliabilism, the scarcity and weakness of explicit arguments for the epistemic closure principle is surprising (Turri 2015a, section 1.2). In defense of the principle, many claim that it is a defining feature of commonsense epistemology. Epistemic closure, they claim, is “just a familiar fact about human knowledge, something we all recognize and abide by in our thought and talk about knowing things” (Stroud 1984, pp. 18-19). The principle deserves to be called “axiomatic” (Cohen 1999, p. 69). Theoretical proposals to abandon closure are “revisionary” and suffer from “deep and wide-ranging” “conflict” with our “intuition” (Hawthorne 2004: 36). It is, we are told, “striking” (Hawthorne 2004: 36) and “startling” (Fumerton 1987) that “some philosophers have gone so far as” to deny the principle and, in the process, “blatantly violate” such a “basic and extremely plausible intuition” (Steup 2005).

For a long time, this extravagant characterization of what “we” find intuitively compelling went unchallenged, even by closure’s opponents (e.g. Nozick 1981, pp. 205-6; Dretske 2013, p. 32). But no serious evidence was ever produced for the claim that closure is a defining feature of

commonsense epistemology. And it is a straightforwardly empirical question whether commonsense exhibits such patterns.

Testing for a commitment to the epistemic closure principle is not difficult. The basic idea is to have people consider simple scenarios and judge whether the agent knows a proposition and one of its obvious, trivial implications. For example, consider this simple scenario (inspired by Vogel 1990, section 2):

When Maxwell arrives at work in the morning, he always parks in one of two spots: C8 or D8. Half the time he parks in C8, and half the time he parks in D8. Today Maxwell parked in C8. It's lunchtime at work. Maxwell and his assistant are up in the archives room searching for a particular document. Maxwell says, "I might have left the document in my car." The assistant asks, "Mr. Maxwell, is your car parked in space C8? It's not unheard of for cars to be stolen." Maxwell thinks carefully for a moment and then responds, "No, my car has not been stolen. It is parked in C8."

Which of the following options best describes Maxwell?

1. He knows that his car is parked in C8. And he knows that his car has not been stolen.
2. He does not know that his car is parked in C8. But he does know that his car has not been stolen.
3. He knows that his car is parked in C8. But he does not know that his car has not been stolen.
4. He does not know that his car is parked in C8. And he does not know that his car

has not been stolen.

If the epistemic closure principle is a defining feature of commonsense epistemology, then when people consider which option best describes Maxwell, the intuitively best answer will not be option 3. Indeed, many philosophers claim that statements like 3 are “abominable,” “ridiculous,” and “distinctly repugnant” (DeRose 1995, pp. 27-29; Dretske 2013, p. 31; Sosa 2004, p. 41). However, when this case was actually tested, nearly two-thirds of people selected option 3 as the best option (Turri 2015a, Experiment 3). That is, the strong central tendency was to attribute *the only* combination of knowledge and ignorance inconsistent with the epistemic closure principle. Nearly no one selected option 3 as best in a closely matched control condition. This same basic pattern persisted across different questioning procedures and narrative contexts (Turri 2015a, Experiments 1, 2, 4, 5; Turri 2015b).

Justification and truth

Suppose that Maria is a watch collector who owns over ten thousand watches. She cannot keep track of all her watches by memory alone, so her accountant maintains a detailed inventory of them. Maria knows that the inventory isn't perfect, but it is extremely accurate. Someone asks Maria, “Does your collection contain a 1990 Rolex?” Maria consults the inventory and it says that she does have one. And this is just another occasion where the inventory is right. Does Maria's evidence justify her in believing that the collection contains a 1990 Rolex? Now imagine a case exactly like the one just described except that it is Mario who collects watches and maintains an inventory, and this is one of those rare cases where the inventory is wrong. Does Mario's

evidence justify him in believing that the collection contains a 1990 Rolex? Can the fact that the inventory is right in the one case but not the other create a difference in what Maria and Mario are justified in believing?

According to the dominant view in contemporary epistemology — the “truth-insensitive” approach to belief evaluation — the answer is a resounding “no.” On this view, Maria and Mario are equally justified because they have “exactly the same reasons for believing exactly the same thing” (Feldman, 2003, p. 29; see also Conee & Feldman, 2004). Their evidence and justification are equivalent across the cases, despite the variation in truth. Focusing on more extravagant examples, consider a staple of contemporary epistemology, the notorious thought experiment involving a radically deceived “brain-in-a-vat.” Leading epistemologists tell us that it is “beyond question at an intuitive level” that such victims are “justified” in their beliefs about their surroundings, just as we are (BonJour, 2003, pp. 185-186). “The intuition” here, we are told, is that you and your victimized counterpart in the super-psychologist’s tank have “equally good evidence” for your respective beliefs (Russell, 2001, p. 38; see also Cohen, 1984, pp. 281-2). We have an “overpowering inclination to think” that the two of you have equally good evidence (Fumerton, 2006, p. 93).

At least one philosopher has claimed that truth profoundly affects what an agent is justified in believing (Sutton, 2007, p. 19; see also Littlejohn, 2013). According to this view, “there are no false justified beliefs” because justification requires knowledge, and knowledge requires truth (Sutton, 2007, p. 1). This view has been labeled an “extraordinary” and “dissident doctrine” lacking in “intuitive credentials” (Conee, 2007). It deviates from “the usual” and “traditional” ap-

proach to questions about belief evaluation, according to which there is no connection between justification and truth (Chisholm, 1989, p. 76).

But is it so extraordinary that truth can directly affect whether a belief is justified? Is it revolutionary to propose that false beliefs are not justified or should not be held? Some philosophers propose that such views are exactly the opposite of extraordinary or revisionary. The extraordinary and revisionist view, they suggest, is the one cultivated in narrow circles within “the professional philosophical community” (Sutton, 2007, p. vii).

So which view is “intuitive” and “natural,” and which is “extraordinary” and “dissident”? Is truth-sensitive epistemology revisionary, or does it accurately reflect our ordinary evaluative practices?

It is easy to test the matter: we just need to elicit judgments about simple cases that vary only the truth-value of the relevant proposition. When this was done, the results showed that ordinary evaluations of belief are deeply sensitive to truth. When confronted with cases very similar to the ones about Maria and Mario above, approximately 75% of people judged that the true belief was justified, but only approximately 25% of people judged that the false belief was justified. The same pattern was observed across different narrative contexts and when using different normative vocabulary to evaluate belief and evidence (e.g. “rational,” “reasonable,” “Maria should believe,” “the responsible thing for Maria to believe is,” “Maria has good evidence to believe,” and “the evidence justifies Maria in believing”) (Turri in press d, Experiment 1; see also Turri in press c, Experiments 1 and 4). Similarly, in a study that asked people to compare the evidence of a normally embodied human and the evidence of his victimized “brain-in-a-vat” twin,

nearly 80% of people judged that the normal human's evidence was better (Turri in press d, Experiment 3).

Knowledge and near-misses

For decades epistemologists have claimed that, intuitively, agents in “fake barn” cases lack knowledge (e.g. Goldman 1976; Sosa 1991: 238-9; Plantinga 1993, p. 33; Pritchard 2005: 161-2; Kvanvig 2008: 274; Pritchard 2014, ch. 6). In a fake barn case, the agent (“Henry” in the original) perceives an object (a barn in the original) and this perceptual relation remains intact throughout. If that is where the story ended, then we would have no problem saying that the agent knows that the object is present. However, it is then revealed, the agent currently inhabits an environment filled with numerous fakes (“papier-mâché facsimiles of barns” in the original), which look like real ones. Epistemologists tell us that, when this information is revealed, we are “strongly inclined” to deny that the agent knows that the object is present. We are then asked, “How is this change in our assessment to be explained?” (Goldman 1976, p. 773).

Any attempt to explain this behavior is idle if the behavior does not occur. And recent research has shown that in fact the behavior does not occur. Instead, people tend to *attribute* knowledge in fake-barn cases (Colaco, Buckwalter, Stich & Machery 2014; Turri in press g). More generally in cases where an agent perceives something that makes her belief true and this perceptual relation remains intact, people overwhelmingly attribute knowledge, even when lookalikes pose a salient threat to the truth of the agent's belief (Turri, Buckwalter & Blouw 2014; Turri in press h). Indeed, commonsense views these cases similarly to paradigm instances

of knowledge, with rates of knowledge attribution often exceeding 80% and not differing from closely matched controls.

Conclusion

Appeals to commonsense and ordinary thought and talk are frequent in philosophy, perhaps nowhere more than in contemporary epistemology. When an epistemological theory implies serious error in commonsense epistemology, it is counted as a cost of the view. It is not necessarily a prohibitive cost, but it is a cost nonetheless. Similarly, when an epistemological theory respects or vindicates deep patterns in commonsense epistemology, it is viewed as a benefit of the view. It is not necessarily a conclusive benefit, but it is a benefit nonetheless.

Standard practice in Anglo-American “analytic” epistemology is to rely on introspection and anecdotal social observation to characterize commonsense epistemology. This method has led to serious mischaracterizations of commonsense and ordinary practice. Above I reviewed five examples of this: knowledge reliabilism, contextualism, the epistemic closure principle, truth-insensitive theories of justification, and knowledge attributions in “fake barn” cases. In each case, the central tendencies in commonsense epistemology are very different from what epistemologists have assumed. The implications of this are potentially significant. I will discuss two.

One implication pertains to how people perceive philosophy. Recently there have been public discussions, in *The New York Times* (Fish 2011) and *The Chronicle of Higher Education* (McIntyre 2011), of whether academic philosophy has any value at all. Extremely successful se-

nior philosophers question the value of philosophy practiced from the armchair (Glymour 2011; Thagard 2012). Many factors probably contribute to negative perceptions of contemporary academic philosophy. One factor is that people are suspicious of attempts to answer philosophical questions from the armchair (Buckwalter & Turri, in press). Indeed, greater prior exposure to academic philosophy is associated with *more* suspicion of armchair philosophical inquiry (Buckwalter & Turri, in press, Experiment 2). A second factor might be that people are alienated by the verdicts that philosophers treat as obvious or commonsensical. For instance, imagine students taking an introductory epistemology course. They are asked to consider a radically deceived “brain-in-a-vat” undergoing experiences qualitatively similar to those of a normal human. “The intuition here,” they are told, is that the brain-in-a-vat and normal human are equally justified in all of their beliefs. This intuition contradicts most people’s judgment about the case, but it is treated as a touchstone for constructing and evaluating theories of justification. The same is true for “the intuition” that the agent in a “fake barn” case does not have knowledge. Before long, some students might conclude that they are not good at philosophy (Buckwalter & Stich 2014, section 4). But, I suspect, it is equally likely that students will conclude, along with Thomas Reid, that if *this* is how philosophers conduct their business, then philosophy is “justly ridiculous” (Reid 1764, p. 21).

Another implication is methodological. To the extent that epistemologists aim to understand the concepts and categories we care about in our ordinary lives — the things which led many of us to study philosophy in the first place — then standard practice in contemporary epistemology should change. As the case studies reviewed above suggest, it is inadequate for epistemologists

to rely only on introspection of their own impressions, the reported impressions of other professionals, and anecdotal social observation. A few simple controlled experiments can quickly cut through mistaken impressions and observations. In short, epistemologists should be quick to rely on controlled experiments to separate wheat from chaff.

This is especially true early in a research program, when results can provide formative clues putting researchers on more promising paths. For example, an enormous amount of time and energy has been devoted to knowledge reliabilism and contextualism over the past few decades. In that time, these views went from novelties to standard topics covered in textbooks, handbooks, anthologies, encyclopedias, and survey courses. Epistemologists lengthily debated finer details of how to understand reliability and the context-sensitivity of knowledge judgments, as well as their relationship to other complicated issues in philosophy. But none of these details matters if the assumptions initiating the discussion are false. Moving forward, some basic preliminary experiments can help to avoid similar false starts. Would that a false start last a few weeks rather than a few decades.

Philosophers are typically not trained to conduct experiments, so this methodological proposal might seem onerous and impractical. But this concern is addressed by two developments. On the one hand, a critical mass of philosophers today can conduct experiments — the experimental philosophers. The number of philosophers fitting this description grows each year. On the other hand, interdisciplinary collaboration is increasingly common, and there have always been trained scientists interested in philosophical issues and willing to work with philosophers on questions of common concern.

Nothing I say here implies the illegitimacy of *prescriptive* epistemology. Epistemologists need not be limited to mapping the contours of commonsense epistemology. Instead, they might wish to cultivate new and better epistemological concepts and categories. But this cannot sensibly be done in ignorance of commonsense concepts and categories. Indeed, quite the opposite — it is silly to prescribe changes unless you understand what you are prescribing changes to! Such understanding is required to intelligently evaluate whether your proposed changes would improve people’s performance. In this respect, experimentation is prerequisite to good prescriptive epistemology.

Of course, not every epistemological project need be experimental and not every epistemologist need be an experimentalist. That would be boring and dysfunctional. It is enough that the field as a whole includes the relevant experimental expertise and brings it to bear in appropriate ways at appropriate junctures. As happens in other fields, researchers will choose a path that suits their interests and temperament — some will focus on more purely theoretical or prescriptive projects, others will focus on more descriptive or experimental projects, others will focus on applied projects, and others will pursue an eclectic mix. The sooner epistemology evolves in this direction, the sooner it will become a more productive, exciting, and respected discipline. Of course, this will probably not happen unless practitioners understand and embrace the value of doing so.

In sum, in light of a string of serious and widespread failures in the field, normalizing experimentation in epistemology is wise. It will help the field avoid unfavorable perceptions and help practitioners identify genuinely commonsensical starting points.

Acknowledgments — For helpful comments and feedback, I thank Wesley Buckwalter, Jennifer Nado, and Angelo Turri. This research was supported by the Social Sciences and Humanities Research Council of Canada and an Early Researcher Award from the Ontario Ministry of Economic Development and Innovation.

References

- Alexander, J., Gonnerman, C., & Waterman, J. (2014). Salience and epistemic egocentrism: an empirical study. In J. R. Beebe (Ed.), *Advances in experimental epistemology* (pp. 97–117). London: Bloomsbury Academic.
- Alston, W. P. (1995). How to think about reliability. *Philosophical Topics*, 23(1), 1–29.
- Aristotle. (350 BCE/1941). *Nicomachean ethics*. (W. D. Ross, Trans., R. McKeon, Ed.).
- Austin, J. L. (1956). A plea for excuses. *Proceedings of the Aristotelian Society*, 57, 1–30.
- BonJour, L., & Sosa, E. (2003). *Epistemic justification: Internalism vs. externalism, foundations vs. virtues*. Malden, Mass.: Blackwell.
- Buckwalter, W. (2010). Knowledge Isn't Closed on Saturday: A Study in Ordinary Language. *Review of Philosophy and Psychology*, 1(3), 395–406. doi:10.1007/s13164-010-0030-3
- Buckwalter, W. (2014). The mystery of stakes and error in ascriber intuitions. In J. R. Beebe (Ed.), *Advances in experimental epistemology* (pp. 145–173). London: Bloomsbury Academic.

- Buckwalter, W., & Stich, S. (2014). Gender and philosophical intuition. In J. Knobe & S. Nichols (Eds.), *Experimental philosophy* (2nd ed.). Oxford: Oxford University press.
- Buckwalter, W., & Turri, J. (in press). Perceived weaknesses of philosophical inquiry: a comparison to psychology. *Philosophia*.
- Chisholm, R. (1989). *Theory of Knowledge* (3rd ed.). Englewood Cliffs, NJ: Prentice Hall.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies*, 46(3), 279–295. doi:10.1007/BF00372907
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57–89.
- Cohen, S. (2013). Contextualism defended. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (2nd ed., pp. 69–75). Malden, Mass.: Wiley-Blackwell.
- Colaco, D., Buckwalter, W., Stich, S., & Machery, E. (2014). Epistemic intuitions in fake-barn thought experiments. *Episteme*, 11(02), 199–212. doi:10.1017/epi.2014.7
- Conee, E. (2007). Review of Jonathan Sutton, *Without Justification*. *Notre Dame Philosophical Reviews*, 12. Retrieved from <http://ndpr.nd.edu/review.cfm?id=11803>
- Conee, E., & Feldman, R. (2004). *Evidentialism: essays in epistemology*. Oxford: Oxford University Press.
- DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review*, 104(1), 1–52.
- DeRose, K. (2009). *The case for contextualism*. Oxford: Oxford University Press.
- DeRose, K. (2011). Contextualism, contrastivism, and X-Phi surveys. *Philosophical Studies*, 156(1), 81–110. doi:10.1007/s11098-011-9799-x

- Dretske, F. I. (1981). *Knowledge and the flow of information*. Cambridge, Mass.: MIT Press.
- Dretske, F. I. (2013). The case against closure. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (2nd ed.). Wiley-Blackwell.
- Feldman, R. (2003). *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Feltz, A., & Zarpentine, C. (2010). Do you know more when it matters less? *Philosophical Psychology*, 23(5), 683–706. doi:10.1080/09515089.2010.514572
- Fish, S. (2011). Does Philosophy Matter? *New York Times*. Retrieved March 2015, from <http://opinionator.blogs.nytimes.com/2011/08/01/does-philosophy-matter>
- Fumerton, R. (1987). Nozick's epistemology. In S. Luper-Foy (Ed.), *The possibility of knowledge*. Rowman & Littlefield.
- Fumerton, R. (2006). *Epistemology*. Malden, MA: Blackwell.
- Glymour, C. (2011). Manifesto. *Choiceandinference.com*. Retrieved March 2015, from <http://choiceandinference.com/2011/12/23/in-light-of-some-recent-discussion-over-at-new-apps-i-bring-you-clark-glymours-manifesto>
- Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(20), 771–791.
- Goldman, A. I. (1993). Epistemic folkways and scientific epistemology. *Philosophical Issues*, 3, 271–285.
- Hansen, N., & Chemla, E. (2013). Experimenting on Contextualism. *Mind & Language*, 28(3), 286–321. doi:10.1111/mila.12019
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.

- Kvanvig, J. (2008). Epistemic Luck. *Philosophy and Phenomenological Research*, 77(1), 272–281. doi:10.1111/j.1933-1592.2008.00187.x
- Lawlor, K. (2013). *Assurance: an Austinian view of knowledge and knowledge claims*. Oxford: Oxford University Press.
- Lemos, N. (2004). *Common sense: a contemporary defense*. Cambridge: Cambridge University Press.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Littlejohn, C. (2013). The Russellian Retreat. *Proceedings of the Aristotelian Society*, 113(3), 293–320. doi:10.1111/j.1467-9264.2013.00356.x
- Locke, J. (1690/1975). *An essay concerning human understanding*. (P. H. Nidditch, Ed.). Oxford: Clarendon Press.
- May, J., Sinnott-Armstrong, W., Hull, J. G., & Zimmerman, A. (2010). Practical Interests, Relevant Alternatives, and Knowledge Attributions: an Empirical Study. *Review of Philosophy and Psychology*, 1(2), 265–273. doi:10.1007/s13164-009-0014-3
- McIntyre, L. (2011). Making Philosophy Matter—or Else. *The Chronicle of Higher Education*. Retrieved March 2015, from <http://chronicle.com/article/Making-Philosophy-Matter-or-130029>
- Nozick, R. (1981). *Philosophical explanations*. Cambridge, Mass: Harvard University Press.
- Plantinga, A. (1993). *Warrant and proper function*. Oxford: Oxford University Press.
- Pritchard, D. (2005). *Epistemic luck*. New York : Oxford University Press.
- Pritchard, D. (2014). *What is this thing called knowledge?* (3rd ed.). New York: Routledge.

- Reid, T. (1764/1997). *An inquiry into the human mind on the principles of common sense*. (D. R. Brookes, Ed.). University Park, Penn.: Pennsylvania State University Press.
- Russell, B. (2001). Epistemic and moral duty. In M. Steup (Ed.), *Knowledge, truth, and duty: essays on epistemic justification, responsibility, and virtue* (pp. 34–48). Oxford: Oxford University Press.
- Sellars, W. (1963). *Science, perception and reality*. Atascadero, CA: Ridgeview Publishing Company.
- Sosa, E. (1991). *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Sosa, E. (2004). Relevant alternatives, contextualism included. *Philosophical Studies*, 119(1), 35–65.
- Steup, M. (2005). Epistemology. *Stanford Encyclopedia of Philosophy*. Retrieved from <http://plato.stanford.edu/entries/epistemology/>
- Stroud, B. (1984). *The significance of philosophical skepticism*. Oxford: Clarendon Press.
- Sutton, J. (2007). *Without Justification*. MIT Press.
- Thagard, P. (2012). Eleven Dogmas of Analytic Philosophy. *Psychologytoday.com*. Retrieved March 2015, from <http://www.psychologytoday.com/blog/hot-thought/201212/eleven-dogmas-analytic-philosophy>
- Turri, J. (2015a). An open and shut case: epistemic closure in the manifest image. *Philosophers' Imprint*, 15(2), 1–18.
- Turri, J. (2015b). Skeptical Appeal: The Source-Content Bias. *Cognitive Science*, 39(2). doi: 10.1111/cogs.12153

- Turri, J. (in press a). Unreliable knowledge. *Philosophy and Phenomenological Research*. doi: 10.1111/phpr.12064
- Turri, J. (in press b). From virtue epistemology to abilism: theoretical and empirical developments. In C. B. Miller, M. R. Furr, A. Knobel, & W. Fleeson (Eds.), *Character: new directions from philosophy, psychology, and theology*. Oxford University Press.
- Turri, J. (in press c). Evidence of factive norms of belief and decision. *Synthese*.
- Turri, J. (in press d). The radicalism of truth-insensitive epistemology: truth's profound effect on the evaluation of belief. *Philosophy and Phenomenological Research*.
- Turri, J. (in press e). A new paradigm for epistemology: from reliabilism to abilism. *Ergo*.
- Turri, J. (in press f). Epistemic contextualism: an idle hypothesis. *Australasian journal of philosophy*.
- Turri, J. (in press g). Knowledge and assertion in "Gettier" cases. *Philosophical psychology*.
- Turri, J. (in press h). Vision, knowledge, and assertion. *Consciousness and cognition*.
- Turri, J., Buckwalter, W., & Blouw, P. (2014). Knowledge and luck. *Psychonomic Bulletin & Review*. doi:10.3758/s13423-014-0683-5
- Vogel, J. (1990). Are there Counterexamples to the Closure Principle? In M. D. Roth & G. Ross (Eds.), *Philosophical Studies Series* (Vol. 48, pp. 13–27). Springer. doi: 10.1007/978-94-009-1942-6_2