THE PROBLEM OF ESEE KNOWLEDGE

JOHN TURRI

University of Waterloo

Traditionally it has been thought that the moral valence of a proposition is, strictly speaking, irrelevant to whether someone knows that the proposition is true, and thus irrelevant to the truth-value of a knowledge ascription. On this view, it is no easier to know, for example, that a bad thing will happen than that a good thing will happen (other things being equal). But a series of very surprising recent experiments suggest that this is actually not how we view knowledge. On the contrary, people are much more willing to ascribe knowledge of a bad outcome. This is known as the *epistemic side-effect effect* (ESEE) and is a specific instance of a widely documented phenomenon, the *side-effect effect* (a.k.a. “the Knobe effect”), which is the most famous finding in experimental philosophy. In this paper, I report a new series of five experiments on ESEE and in the process accomplish three things. First, I confirm earlier findings on the effect. Second, I show that the effect is virtually unlimited. Third, I introduce a new technique for detecting the effect, which potentially enhances its theoretical significance. In particular, my findings make it more likely, though they do not entail, that the effect genuinely reflects the way we think about and ascribe knowledge, rather than being the result of a performance error.

1. Introduction

Most theorists think that whether a belief counts as knowledge is unaffected by the desirability of the belief or its content. If we are wondering whether you know that a certain event will occur, we will attend to whether you believe that the event will occur, whether it is true that the event will occur, whether you have any evidence that it will occur, and perhaps also whether your evidence only “luckily” led you to believe the truth. It does not matter whether the event in question is good or bad. That is, other things being equal, the event’s valence does not affect whether you know that it will occur. Or so mainstream theorizing about knowledge would

---

Contact: John Turri <john.turri@gmail.com>

have us believe (e.g. Steup 1996; Lehrer 2000; Feldman 2003; Fumerton 2006; Bonjour 2009).

But important recent experimental work suggests that mainstream opinion is simply wrong about this (Beebe and Buckwalter 2010; Beebe and Jensen 2012; Beebe and Shea forthcoming). The valence of an outcome apparently does influence whether we ascribe knowledge. For example, people are more willing to ascribe knowledge to a CEO if he thinks that his company's policies will *harm* rather than *help* the environment. This is known as the *epistemic side-effect effect* (ESEE; pronounced *easy*). It is a specific instance of a widely documented phenomenon, the *side-effect effect* (a.k.a. "the Knobe effect"), which is the most famous finding in experimental philosophy. ESEE is an important and surprising discovery about how we think of knowledge which should be of interest to philosophers and cognitive scientists more generally. It is my focus here.

Contemporary epistemologists have taken it for granted that patterns in ordinary usage should, or at least can, constrain substantive epistemological theorizing. It is widely assumed that theorists should respect patterns of competent and literal knowledge attribution. Other things being equal, we prefer an epistemological theory to the extent that it counts ordinary knowledge attributions as competent and literal (more on this in section 2.2). In this paper, I investigate whether we should think that ESEE reflects such a pattern and, in turn, whether it ought to inform epistemological theory. Although further work ultimately remains to be done, I report some good initial evidence that ESEE does reflect a pattern worthy of theoretical respect.

Here is the paper's plan. Section 2 situates the present discussion within a broader research program; it reviews previous work on side-effect effects and outlines competing approaches to explaining them. Section 3 poses two questions that set the stage for further experimental inquiry into ESEE. Section 4 reports four experiments which collectively suggest that, when observed using previous methods, ESEE is too broad and too easily activated for it to reveal how we think about knowledge. Section 5 therefore suggests an alternative approach for detecting ESEE. Section 6 enacts this alternative and reports an experiment that provides the best evidence thus far that ESEE reflects a competent pattern of knowledge ascription. Section 7 sets an agenda for further research into this topic.

2. Knowledge Can Be ESEE

There has been a spate of valuable recent work on ordinary knowledge ascriptions (e.g. Weinberg et al. 2001; Swain, Alexander, and Weinberg 2008; Nagel 2008; Feltz and Zarpentine 2010; Starmans and Friedman 2012; Schaffer and Knobe 2012; Sripada and Stanley, 2012). These studies suggest that some surprising fac-

tors affect whether laypeople attribute knowledge to someone (for an overview, see Buckwalter 2012). Here I will focus on one set of studies in particular. These studies take inspiration from Joshua Knobe's pioneering work on how people's evaluative judgments affect whether they describe outcomes as *intentional* (Knobe 2003a, 2003b, 2004, 2010). This section briefly reviews relevant prior work.

2.1. SEE, *This Is Easy*

Suppose that an agent embarks on a course of action aimed at a primary outcome (a central effect). And suppose that the agent also anticipates but does not desire that her conduct will have a certain side effect. Knobe found that if the side effect is bad or undesirable, then people are much more likely to describe the agent as having *intentionally* brought it about. For example, to use Knobe's now famous case, suppose that a CEO decides to start a new program with the primary aim being to increase profits. And suppose that the CEO also anticipates but does not desire that the program will, as a side effect, affect the environment. Did the CEO intentionally bring about the environmental side effect? People are much more likely to say that he did if the side effect is *harmful* rather than *helpful*. This is known as *the side-effect effect* (SEE).

A side effect's valence influences ascriptions of intentionality, and it has also been shown to affect ascriptions of other practical attitudes, such as *deciding*, *being in favor of*, and *desiring*, among others (Pettit and Knobe 2009), as well as *act individuation* (Ulatowski 2009). This is true whether the side effect is negatively valenced because it violates a norm of morality, prudence, aesthetics, or etiquette. The effect appears to be present regardless of gender, age, or native language. The effect is very robust.

An important question is what explains SEE. Knobe groups the competing accounts into *competence explanations* and *distortion explanations*, and this taxonomy has taken hold in the literature (e.g. Buckwalter 2013). Competence explanations attribute the effect to the fact that "moral considerations actually figure in the competencies people use to" attribute psychological states (Knobe 2010: 320). Knobe favors this sort of account, though he does not claim that it has any immediate implications for a substantive theory of what psychological states are, or for a semantic theory of psychological predicates (e.g. Knobe and Burra 2006: 332–333, 338). Instead, Knobe claims only that the effect reveals something deep and important about the way we ordinarily think about psychological states, about our underlying competence in applying such concepts.

Distortion explanations deny that moral considerations figure into the underlying competence and instead attribute the effect to other factors that distort the competence's operation. They claim that the effect is due to performance error (Nadelhoffer 2006; Malle 2006: 103–104; Alicke and Rose 2010) or to pragmatic

concerns that do not reveal anything about how people think about psychological states per se (Adams and Steadman 2004) or to features of the materials used in the experiments (Guglielmo and Malle 2010).

Knobe gently sets aside whether there is a semantic or metaphysical upshot of his conceptual competence model. But due to certain preoccupations and methodological commitments in contemporary epistemology, when we turn to whether the side-effect effect shows up for knowledge ascriptions too, it is not so simple to maintain the division between people's underlying competence in classifying things in a certain category, on the one hand, and the nature of things that populate that category, on the other.¹

Epistemologists have taken it for granted that patterns in ordinary usage should, or at least can, constrain substantive theorizing about the nature of knowledge or the (abstract) concept of knowledge or the meaning of “knows” (Austin 1956; Cohen 1999; Rysiew 2001; DeRose 2009; Hawthorne 2004; Stanley 2005; Fantl and McGrath 2009). Of course, not just any pattern is ripe for guiding theory. But it is widely assumed that theorists should respect patterns of competent and literal knowledge ascription. There will always be some interplay between one's theory and what one is willing to consider a pattern of competent and literal knowledge ascription (e.g. Davis 2007; Bach 2008; Turri 2014). But it is not “anything goes.” And a theory which implies that much of our ordinary practice of knowledge ascription is either incompetent or non-literal has, as they say, “some explaining to do.”

Setting aside contemporary epistemologists' methodological predilections, there is another reason to take competent and literal application of concepts seriously. Over the past forty years, philosophers of language and mind have developed *externalist* theories of semantic meaning and mental content (Kripke 1972; Putnam 1975; Burge 1979; Burge 1986). Though not uncontroversial, these externalist theories are widely accepted. If, as Donald Davidson (1983: 146) puts it, “Belief is in its nature veridical,” or as George Santayana (1923: 9) wrote, “Intelligence is by its nature veridical,” or as the 11th-century Indian philosopher Sriharsa thought, “Cognition is ordinarily by nature true or veridical” (Phillips 2011), then we should expect competent and literal applications of concepts to be largely accurate. So despite Knobe's attempt to downplay it, a conceptual-competence account might have a fairly direct semantic or metaphysical upshot after all.

In what follows, my focus will be knowledge, knowledge ascriptions and the concept of knowledge. I do not claim that any point I make generalizes to other psychological states, ascriptions, concepts, or side-effect effects concerning them.

---

1. Buckwalter (2013) does an especially good job of explaining and properly emphasizing the significance of these methodological points.

2.2. *See, This Is ESEE*

James Beebe and Wesley Buckwalter (2010) showed that the side-effect effect extends to knowledge ascriptions too. They took Knobe's original case and instead of asking participants whether the CEO intentionally brought about the environmental side effect, they asked whether the CEO knew that the program would bring about the side effect. Beebe and Buckwalter observed a striking asymmetry: Participants were significantly more likely to ascribe knowledge if the environment is harmed rather than helped.² In other words, they discovered ESEE. Their result has been replicated and appears to be robust (Beebe and Jensen 2012; see also section 4.1 below).

Beebe and Buckwalter conjecture that ESEE might reveal something important about the nature of knowledge and its relation to moral value: Whether a belief counts as knowledge can depend on the moral status of actions based on it. But they are careful to note that there is no easy path from the observed patterns of knowledge attribution among laypeople to sound theoretical conclusions about the nature of knowledge. Beebe and Jensen (2012) are likewise cautious about inferring anything about the nature of knowledge itself from the ESEE data. But Beebe, Buckwalter and Jensen all seem to agree that the ESEE data lend support to a view about the nature of knowledge which has lately been gathering support, namely, that factors other than truth, belief, evidence and reliability can directly affect whether you know something (e.g. Fantl and McGrath 2002, 2009; Hawthorne 2004; Stanley 2005). On this approach, practical, moral and other traditionally *nonepistemic* evaluative factors can directly affect whether you know something. This is an intriguing thesis which, if true, reveals something deep and important about knowledge.

2.3. *Gee, This Is ESEE*

Buckwalter (2013) and Beebe and Shea (forthcoming) have recently extended this line of reasoning to "Gettier cases." Very roughly, a Gettier case is a case where a subject has a justified true belief which, according to conventional wisdom in philosophy, fails to count as knowledge due to an improbable confluence of luck (Gettier 1963; Starmans and Friedman 2012; Turri 2011, 2012a, 2013). In particular, Gettier cases involve two strokes of luck: a stroke of bad luck, which would ordinarily prevent the subject's belief from being true, followed by a stroke of good luck, which "cancels out" the bad luck by making the subject's belief true anyway (Zagzebski's 1994: 66, 1996: 288–289; compare Sosa 1991: 238).

---

2. Right around the same time, some armchair epistemologist got lucky and predicted the same thing (see Turri 2012a).

Beebe and Shea asked participants about a CEO very similar to the one in Knobe's original story, except that this CEO has been "Gettiered." Participants were asked to agree or disagree with the claim that the CEO knew that the environment would be helped or harmed. Participants rated their agreement or disagreement on a 7-point Likert scale. They were significantly more likely to agree with the statement if the side effect was harmful rather than helpful. The mean response in harmful conditions (4.35) fell above the mid-point and responses of 5 and 7 were both modes, indicating that participants tended to think that the Gettiered CEO knew about the effect on the environment. Beebe and Jensen observed similar results with different cover stories. Buckwalter has observed even more striking results with some of his vignettes, with mean responses topping 6 on a 7-point scale. This is *Gettier made ESEE*, or the *Gettier epistemic side-effect effect* (GESEE). I replicated these results (see section 4.2 below).

A natural reaction to these cases is that participant response is being distorted by a desire to blame the protagonist and to hold him accountable for a rotten outcome that he could have prevented (compare Alicke 2008, Nadelhoffer 2006). Agreeing that the CEO "knew" is a way of agreeing that he should be held responsible for the environmental harm. By contrast, people are not interested in giving the CEO credit for environmental improvement that he did not even care about. But Buckwalter (2013) has observed the effect persist even when participants are asked whether some *third party*, who cannot plausibly be held responsible, knows that the environmental side effect will occur. So a straightforward blame-based distortion account does not explain all the data.

If we are attracted to the idea that our ordinary practice of ascribing knowledge should at least broadly constrain theorizing about knowledge itself, then these data seriously threaten some well-entrenched positions in contemporary epistemology. Perhaps whether you know does partly depend on whether you are up to no good. Perhaps we should leave room in our epistemology for Gettiered knowledge.

3. How Easy Can It Get?

Maybe ESEE can be extended to unsettle other articles of conventional wisdom too. Conventional wisdom has it that knowledge requires truth and that knowledge requires belief. Are these requirements reflected in ordinary usage?

3.1. *It's Not Easy Being . . . CHESEE?*

Consider this contrarian hypothesis:

Contrarian Hypothesis (CH): Knowledge is not factive.

Otherwise put, CH says that knowledge does not require truth: It is possible to know false things. Some philosophers have argued that, on the ordinary conception of knowledge, CH is true (Hazlett 2010, 2012). But virtually all philosophers think that CH is obviously false. And Buckwalter (ms.) has reported several studies which suggest that, on the ordinary conception of knowledge, CH is false. But Buckwalter cautiously reminds us that there might be other evidence that knowledge, ordinarily understood, is not factive. We saw SEE extended to ESEE, which was in turn extended to GESEE. Perhaps the same phenomenon might also provide evidence for CH. In other words, we have seen that knowledge ascriptions can be made ESEE, so can they also be made CHESEE (pronounced *cheesy*)?

3.2. *Is Knowledge a Breeze?*

Consider also this unorthodox view:

Belief not Required (BR): Knowledge does not require belief.

Some philosophers have advocated BR (e.g. Radford 1966; Lewis 1996), and experimental philosophers have recently taken up its cause (Schwitzgebel and Myers-Schulz forthcoming; Murray, Sytsma and Livengood 2012; Buckwalter, Rose and Turri 2013). But most philosophers think that BR is obviously false; they think it is *obvious* that knowledge entails belief (Rose and Schaffer 2013). But perhaps that is not the ordinary view. Perhaps the ESEE pattern will extend to BR too. We have seen that knowledge ascriptions can be made ESEE. Can they also be made BRESEE (pronounced *breezy*)?

4. Actually, This Is too Easy

Suppose that in the same way that knowledge ascriptions can be made ESEE and GESEE, they can also be made CHESEE or BRESEE or both. What would that show?

I think that it would make a conceptual-competence explanation of the observed effect much less likely. It would seem more likely that participants are incompetently applying their concept of knowledge, or competently but falsely applying it for pragmatic reasons, or competently applying some other concept in response to a “knowledge” question. Over thousands of years and across many different cultures, careful reflection has led people repeatedly to the view that knowledge requires truth and belief (or something very similar). This is overwhelmingly reflected in the most influential historical and contemporary theories of knowledge (Matilal 1986; Steup 1996; Lehrer 2000; Feldman 2003; Fumerton 2006; Bonjour

2009; Phillips 2011). And if a conceptual-competence explanation is unlikely, then we need to be much more cautious when determining whether the ESEE data support various substantive proposals about the nature of knowledge itself. Of course, none of that would go to show that the effect is unreal. On the contrary, the point is that the effect is, ironically, *too real* for it to guide a theory of knowledge.

By contrast, if knowledge ascriptions cannot be made CHESEE or BRESEE in that same way, it eliminates two barriers to a conceptual-competence account. Of course, it would not entail that a conceptual-competence account is true. But it would eliminate potential worries about the effect not ruled out by previous studies.

4.1. Experiment 1

This section reports an experiment designed to test whether knowledge ascriptions can be made CHESEE. That is, the experiment tests whether the method used to detect ESEE and GESEE in earlier experiments will also produce results that support the unorthodox view that knowledge does not require truth.

4.1.1. Method

Participants ($N = 147$)³ were randomly assigned to one of six conditions in a 2 (Valence of side-effect: Help/Harm) $\times$ 3 (Luck: Normal/False/Gettier) between-subjects experiment: Normal Harm, False Harm, Gettier Harm, Normal Help, False Help, Gettier Help. (See Table 1.) Each participant read one of six stories.

The stories for the conditions were variants of the CEO case, built up from narrative modules. These were the narrative modules (help/harm variations bracketed and separated by a slash):

(*No Luck*) The vice president of a manufacturing company went to the CEO and said, “We are thinking of starting a new program. It will increase our profits, [and/but] it will also [improve/harm] local water quality, because it requires that we [start/stop] dumping polluted water into the river behind

---

3. Sixty-seven female, aged 18–65, $M = 28.8$, $SD = 9.92$. As with the experiments reported below, participants were recruited using Amazon Mechanical Turk and compensated \$0.25 for approximately 2–3 minutes of their time. Participants were not allowed to re-take any survey reported here, and participants who had taken previous similar surveys were excluded by their AMT Worker ID. Participants were located throughout the United States. Ninety-eight percent reported English as their native language. They filled out a brief demographic survey after testing. I excluded data from fifty-five participants who failed comprehension questions. Including data from these participants made a small but statistically insignificant difference to the results reported below. Previous studies of ESEE effects do not report asking comprehension questions.

our manufacturing plant.” The CEO answered, “I don’t care at all about [improving/harming] local water quality. I just want to make as much profit as I can.” Then the CEO made his decision: “Let’s start the new program and make some serious profits.” The company then started the new program, and sure enough, over the next few months, it caused their profits to increase dramatically.

(*Bad Luck*) But as it turns out, the vice president was wrong about something: the company’s new program did not require them to [start/stop] dumping polluted water into the river. So they never did [start/stop].

(*Good Luck – Harm*) However, at the same time that the company was starting its new program, a local military installation decided to save money on expensive waste disposal. Instead of going through the proper procedures, they decided to secretly dump all of their toxic waste straight into a local lake, which harmed local water quality. This illegal and harmful dumping remained a secret and was never made public. When the CEO made his decision to start the new program, he was unaware that the military was going to dump the toxic waste.

(*Good Luck – Help*) However, at the same time that the company was starting its new program, the mayor decided to take action and improve local water quality. Instead of going through the proper procedure, the mayor secretly ordered the local water department to upgrade the filtration system at the local water treatment plant, which improved local water quality. This illegal but helpful expenditure remained a secret and was never made public. When the CEO made his decision to start the new program, he was unaware that the mayor was going to order the upgrade.

Table 1 shows how stories for the different conditions were built, along with the questions used. The “Good Luck” and “Bad Luck” modules combine to form Gettier cases by virtue of the “double luck” recipe described earlier (section 2.3). These were the comprehension questions (options in brackets):⁴

(CQ₁) Did the company’s program cause their profits to increase? [Yes/No]

(CQ₂) Did the company’s program improve/harm local water quality? [Yes/No]

(CQ₃) When the CEO makes his decision to start the program, he thinks that local water quality will be [improved/harmed] because he thinks that _____. [the CEO’s company will start/stop dumping pollution into the river/the mili-

---

4. Options were rotated randomly for all questions in all experiments reported here, except for Likert scales and confidence measures, which were always ordered low-to-high. Likewise, questions were always presented in the same order in all experiments.

tary will start dumping toxic waste into the lake/the local water treatment plant will be upgraded]
(CQ4) What causes local water quality to be improved/harmed? [The CEO’s company started/stopped dumping pollution into the river./The military started dumping toxic waste into the lake./The local water treatment plant was upgraded.]⁵

The test question was:

(Klikert) Please tell us the extent to which you agree or disagree with the following statement: At the point in the story when the CEO makes his decision to start the program, he knows that local water quality will be harmed.

Responses were collected on a five-point scale: strongly disagree (= 1), disagree, neutral, agree, strongly agree (= 5). Questions were all presented on a single screen and the story remained at the top of the screen throughout. Questions were always presented in the same order. Response options for the questions were rotated randomly, except for the options on Likert scales, which were always presented in the same order.⁶

TABLE 1. Experiment 1: List of modules composing the storylines, along with the questions, in the different conditions.

Normal Harm	False Harm	Gettier Harm
No Luck	No Luck	No Luck
CQ1–4	Bad Luck	Bad Luck
Klikert	CQ1–3	Good Luck – Harm
	Klikert	CQ1–4
		Klikert
Normal Help	False Help	Gettier Help
No Luck	No Luck	No Luck
CQ1–4	Bad Luck	Bad Luck
Klikert	CQ1–3	Good Luck – Help
	Klikert	CQ1–4
		Klikert

5. CQ4 was not asked in the False Harm and False Help conditions because it confusingly presupposes that local water quality was improved/harmed. The stories for these conditions give no indication that anything does improve/harm local water quality.

6. The same is true in all experiments reported here.

Setting up the experiment this way allowed me to accomplish three things at once. First, comparing the Normal conditions allows us to test whether we can replicate earlier experiments in which participants were more likely to count a true belief as knowledge in the Harm condition than in the Help condition; that is, it tests ESEE. Second, comparing the Gettier conditions allows us to test whether we can replicate earlier experiments in which participants were more likely to count a Gettiered belief as knowledge in the Harm condition than in the Help condition; that is, it tests GESEE. Third, comparing the False conditions allows us to test whether participants are more likely to count a false belief as knowledge in the Harm condition than in the Help condition; that is, it tests CHESEE.

4.1.2. Results and Discussion

A one-way analysis of variance (ANOVA) revealed a main effect of Valence⁷ and an interaction effect between Valence and Luck.⁸ There was no main effect of Luck.⁹

The results from the Normal conditions replicated earlier findings on ESEE. Knowledge scores in Normal Harm were higher than in Normal Help.¹⁰ The mean knowledge score in both conditions was significantly above the midpoint.¹¹ The mode in Normal Harm was 5, compared to a mode of 4 in Normal Help.

The results from the Gettier conditions replicated earlier findings on GESEE. Knowledge scores in Gettier Harm were trending higher than in Gettier Help.¹² The mean knowledge score in both conditions was significantly above the midpoint.¹³ The mode in Gettier Harm was 5, compared to a mode of 4 in Gettier Help.

The results from the False conditions provide equally good evidence for CHESEE. Knowledge scores in False Harm were higher than in False Help.¹⁴ The mean knowledge score in the Harm condition was significantly above the midpoint¹⁵ and virtually identical to Normal Harm. The mode in False Harm was 5, compared to a mode of 4 in False Help.

Figure 1 visualizes the results from this experiment.

7. $F(1, 141) = 34.9, p < .001, \eta_p^2 = .198$, all tests two-tailed unless otherwise noted.

8. $F(2, 141) = 5.2, p = .013, \eta_p^2 = .059$.

9. $F(2, 141) = 1.48, p = .231$, n.s.

10. Normal Harm, $M = 4.79, SD = 0.77$; Normal Help, $M = 3.77, SD = 1.24$; $t(40) = 3.28$, $MD = 1.02, p = .002$.

11. Normal Harm, $t(28) = 12.5, p < 0.001$; Normal Help, $t(12) = 2.245, p = 0.044$; test value = 3.

12. Gettier Harm, $M = 4.24, SD = 1.23$; Gettier Help, $M = 3.74, SD = 1.15$; $t(54) = 1.56$, $MD = 0.5, p = .0625$, one-tailed.

13. Gettier Harm, $t(24) = 5.02, p < .001$; Gettier Help, $t(30) = 3.6, p = .001$; test value = 3.

14. False Harm, $M = 4.76, SD = 0.83$; False Help, $M = 3.0, SD = 1.25$; $t(39.76) = 5.824$, $MD = 1.76, p < .001$.

15. $t(24) = 10.6, p < 0.001$; test value = 3.

Figure 1: Experiment 1: Mean agreement with the knowledge ascription, on a 1 (strongly disagree) to 5 (strongly agree) scale.

4.2. Experiment 2

This section reports an experiment designed to test whether knowledge ascriptions can be made BRESEE. That is, the experiment tests whether the method used to detect ESEE, GESEE and CHESEE in earlier experiments will also produce results that support the unorthodox view that knowledge does not require belief.

4.2.1. Method

Participants ($N = 56$)¹⁶ were randomly assigned to one of two conditions: Harm and Help. Participants read a minor variant of a CEO case, in which the CEO starts out by saying, “I agree that this plan will make us lots of profit. But I reject your claim that it will improve/harm local water quality.” Here is the full vignette (help/harm variations bracketed and separated by a slash):

The vice president of a manufacturing company went to the CEO and said,
“We are thinking of starting a new program. It will increase our profits,

16. Twenty female, aged 18–62, $M = 28.5$, $SD = 10.75$. I eliminated data from 14 participants who failed comprehension questions. Ninety-nine percent reported English as a native language.

[and/but] it will also [improve/harm] local water quality, because it requires that we [start/stop] dumping polluted water into the river behind our manufacturing plant.” The CEO answered, “I agree that this plan will make us lots of profit. But I reject your claim that it will [improve/harm] local water quality. Anyway, I don’t care at all about [improving/harming] local water quality. I just want to make as much profit as I can.” Then the CEO made his decision: “Let’s start the new program and make some serious profits.” ¶¹⁷ The company then started the new program, and sure enough, over the next few months, they made lots of profits. But the vice-president was right: local water quality was also [improved/harmed].

The comprehension questions were:

- (CQ1) Did the company’s program cause water quality to be harmed/improved? [Yes/No]
 (CQ2) When the CEO makes his decision to start the new program, he _____ the vice president’s claim that it will improve/harm local water quality. [rejects/accepts]

The test question was the same as in Experiment 1.

4.2.2. Results and Discussion

The CEO in each version of the case rejects the claim that the program will harm local water quality. If the earlier pattern of ESEE results holds, then there will be an effect of condition, with Knowledge scores higher in the Harm condition.

This is exactly what was observed. Knowledge scores were higher in the Harm condition than in the Help condition.¹⁸ The mean knowledge score was significantly above the midpoint in Harm¹⁹ and significantly below the midpoint in Help.²⁰ The modes in Harm were 4 and 5, compared to a mode of 2 in Help.

It might be thought that participants in Harm were attributing to the CEO belief in the relevant claim. But this seems unlikely because there was a comprehension question (CQ2) to eliminate such participants. Only participants who answered that the CEO rejects the claim were included in the analysis. Moreover, there was no effect of condition on whether participants failed CQ2.²¹ Figure 2 (section 4.4.2 below) visualizes the results from this experiment.

17. Indicates paragraph break as it appeared on the participant’s screen.

18. Harm, $M = 3.6$, $SD = 1.36$; Help, $M = 2$, $SD = 1.05$; $t(50.38) = 4.64$, $MD = 1.6$, $p < .001$.

19. $t(34) = 2.62$, $p = .013$.

20. $t(20) = -4.37$, $p < .001$.

21. Fisher’s exact test, $p = .517$

4.3. *Experiment 3*

We have seen that participants agree that the CEO knows when the proposition in question is false, and when the CEO does not believe the proposition. But suppose that the proposition is both false and not believed. Will participants then disagree with the knowledge ascription? This section reports an experiment that answers that question. (Unfortunately, I was unable to invent a memorable acronym for this effect.)

4.3.1. Method

Participants ($N = 57$)²² were randomly assigned to one of two conditions: Harm and Help. Participants in this experiment received the exact same treatment as their counterparts in Experiment 2, except that the final sentence of the story was, “But the vice president was wrong about something: local water quality was not improved/harmed.”

4.3.2. Results and Discussion

Once again, knowledge scores were higher in Harm than in Help.²³ The mean knowledge score did not differ significantly from midpoint in Harm,²⁴ whereas it was significantly below the midpoint in Help.²⁵ The modes in Harm were 2 and 4, compared to a mode of 1 in Help.

4.4. *Experiment 4*

I went even further in the quest to get participants to disagree with the knowledge ascription. I presented them with stories that take away not only truth or belief, but also justification and, in one case, *all three*: belief, truth and justification. I dispensed with the comparison to a Help condition at this point. Would participants finally disagree?

---

22. Twenty-nine female, aged 18–71, $M = 26.9$, $SD = 8.84$. One hundred percent reported English as a native language. I excluded data from fifteen participants who failed comprehension questions.

23. Harm, $M = 2.94$, $SD = 1.21$; Help, $M = 1.73$, $SD = 1.28$; $t(55) = 3.6$, $MD = 1.21$, $p < .001$.

24. $t(34) = -0.279$, $p = .782$, n.s., test value = 3.

25. $t(21) = -4.67$, $p < .001$, test value = 3.

4.4.1. Method

Participants ($N = 47$)²⁶ were randomly assigned to one of two conditions: True and False. Each participant read one of two stories. In each story, the CEO is not justified in believing that the program will harm local water quality. The vice president says,

We are thinking of starting a new program. It will increase our profits, but there is a very, very small chance that it will harm local water quality, since it is just possible that it will require us to dump polluted water into the river behind our manufacturing plant. Nevertheless, that almost certainly won't happen.

The CEO then explicitly rejects the claim that *it is even possible* that the program will harm water quality. Participants in the True condition read a story in which it nevertheless turns out that, “against all odds,” the program harms water quality. Participants in the False condition read a version in which there “was never a real chance” that the program would harm water quality. (The full vignette is included in an Appendix.) The comprehension and test questions were the same as in Experiment 2.

4.4.2. Results and Discussion

A one-way ANOVA revealed no effect of condition on response to the knowledge attribution.²⁷ In neither condition did mean response differ significantly from the midpoint.²⁸ The mode in False was 2 (followed closely by 5!), while 2 and 5 were both modes in True.

In the story for the False condition in this experiment, the CEO is not justified in believing the claim in question, he rejects the claim, and the claim is false. But participants still were *overall neutral* on whether this *unjustified false non-belief* is knowledge! Perhaps if they were told that the CEO had died the week before the proposed program was even invented, then they would unambiguously disagree that he knows. Figure 2 visualizes the results from Experiments 2–4.

26. Twenty-three female, aged 19–65, $M = 29$, $SD = 10.2$. Ninety-five percent reported English as a native language. I eliminated data from fourteen participants who failed comprehension questions.

27. $F(1, 45) = 0.714$, $p = .403$, n.s.

28. True, $M = 3.25$, $SD = 1.54$, $t(23) = 0.796$, $p = .434$, n.s.; False, $M = 2.87$, $SD = 1.55$, $t(22) = -0.13$, $p = .690$, n.s.; test value = 3.

Figure 2: Experiments 2–4: Mean agreement with the knowledge ascription, on a 1 (strongly disagree) to 5 (strongly agree) scale. Experiment 2: True non-belief; Experiment 3: False non-belief; Experiment 4: Unjustified true/false non-belief. No Help condition was included in Experiments 3 and 4.

5. Re-evaluating

Experiments 1–3 show that participants are consistently more willing to agree that the CEO knows that the environmental side effect will occur when it is harmful rather than helpful. This is true whether the CEO has a normal justified true belief, a Gettiered belief, a false belief, no belief, and even when the CEO has no belief and the proposition in question is false. Moreover, participant agreement with the knowledge ascription was very high in the case of a false belief. Participants in Harm conditions were neutral on whether false non-beliefs were knowledge, on whether unjustified true non-beliefs were knowledge, and even on whether unjustified false non-beliefs were knowledge.

If we take these results at face value, then there is virtually nothing we can do to make participants *disagree* that the CEO knows that the harmful side effect will occur. This poses a serious challenge to the conceptual-competence account of epistemic side-effect effects, for reasons noted at the beginning of section 4. It seems unlikely that participants in these studies are competently and literally applying the concept of knowledge. A more likely explanation is that a seriously negative reaction to the CEO is causing performance errors, or that at least many participants are agreeing to a statement other than the one explicitly featured in the

test question, perhaps along the lines of, “The CEO is a world-class jerk.”²⁹ It also seems unlikely that participants are competently and literally applying the concept of *confidently held belief*, since they persist in agreeing that the CEO knows even when participants explicitly acknowledge that the CEO rejects the claim.

6. The Big ESEE: A Different Approach

I am interested in trying a different approach to detecting epistemic side-effect effects. It would be good to detect an effect in such a way that it stands a fighting chance of revealing something about our conceptual competence in ascribing knowledge, and perhaps, in turn, about knowledge itself. This requires not only detecting an effect, but an *appropriately circumscribed* effect.

For this to work, participants must also be willing to *deny* knowledge in a certain range of control cases when probed in the same way that produces ESEE. For example, they must be willing to deny that false beliefs are knowledge. More generally, it would help—though it is *not* necessarily required—if we detected ESEE in an overall pattern of knowledge attribution that broadly agrees with mainstream theorizing about knowledge. The agreement does not have to be perfect. But if the two are totally at odds, then many epistemologists will suspect that the observed effect is not due to conceptual competence but rather to performance error or more pressing practical concerns, such as blaming the CEO or expressing disapproval.

This section reports an experiment designed to detect ESEE within such a pattern. The experiment takes a different approach to questioning participants. It features a binary knowledge question whose options are “really knows” and “only thinks he knows,” along with a confidence measure. This basic approach, pioneered by Christina Starmans and Ori Friedman (2012), has proven effective in the past (see also Turri 2013; but see Cullen 2010 for some cautionary points). It is worth trying here.

6.1. Experiment 5

6.1.1. Method

Participants ($N = 228$)³⁰ were randomly assigned to one of six conditions (Table 2) in the same 2 (Valence: Help/Harm) $\times$ 3 (Luck: Normal/False/Gettier) between-subjects design as in Experiment 1. (See Table 2.) The stories for the conditions were again built up from narrative modules. Table 2 shows how the stories were

29. Compare Alicke and Rose 2010.

30. Eighty-three female, aged 18–67, $M = 27.2$, $SD = 9.75$. Ninety-seven percent reported English as a native language.

built, along with the questions used. The narrative modules were the same as in Experiment 1, as were the comprehension questions. But the test question was different:

(KQ) When the CEO makes his decision to start the new program, he _____ that local water quality will be improved/harmed [really knows/only thinks he knows].

Participants were then asked to rate how confident they were in their answer to the test question, 1 (not at all confident) to 10 (completely confident).

TABLE 2. Experiment 5: List of modules composing the storylines, along with the questions.

Normal Harm	False Harm	Gettier Harm
No Luck	No Luck	No Luck
CQ1–4	Bad Luck	Bad Luck
KQ	CQ1–2, 4	Good Luck
	KQ	CQ1–4
		KQ
Normal Help	False Help	Gettier Help
No Luck	No Luck	No Luck
CQ1–4	Bad Luck	Bad Luck
KQ	CQ1–2, 4	Good Luck
	KQ	CQ1–4
		KQ

As before, the Normal stories are very similar to the original CEO cases. The False stories introduce a twist: it turns out that the environmental side effect does not occur. The Gettier stories introduce *two* twists: the new program does not produce the environmental side effect, but the something else does, which the CEO is unaware of.

6.1.2. Results and Discussion

Following Starmans and Friedman (2012), I define a *weighted knowledge ascription* as the product of the answer to the dichotomous knowledge question (really knows = 1; only thinks he knows = -1) and the reported confidence (1–10). Scores for this measure thus fall on a twenty-point scale, ranging from -10 (fully confident knowledge denial) to 10 (fully confident knowledge ascription).

A one-way analysis of variance (ANOVA) detected a main effect of Valence,³¹

31. $F(2, 222) = 37.26, p < .001, \eta_p^2 = .251$.

a main effect of Luck,³² and an interaction of Valence and Luck³³ on weighted knowledge ascription.

Pairwise comparisons with independent-samples t-tests revealed that mean weighted knowledge ascription was higher in Normal Harm than in Normal Help,³⁴ higher in False Harm than in False Help,³⁵ but no different between Gettier Harm and Gettier Help.³⁶ Mean weighted knowledge ascription was significantly above the midpoint in Normal Harm,³⁷ no different from midpoint in Normal Help,³⁸ and significantly below the midpoint in both False conditions and both Gettier conditions.³⁹

Looking at the dichotomous knowledge question reveals a similar picture. The rate of knowledge attribution was higher in Normal Harm than in Normal Help,⁴⁰ and it was higher in False Harm than in False Help,⁴¹ but it did not differ between Gettier Harm and Gettier Help.⁴² Rate of knowledge attribution was above what could be expected by chance in Normal Harm,⁴³ no different from chance in Normal Help,⁴⁴ and significantly below chance in all the other conditions.⁴⁵

Figure 3 visualizes the results from this experiment:

These results provide evidence of an appropriately constrained epistemic side-effect effect. Overall, only one condition produced results that are arguably inconsistent with a mainstream, nonskeptical theory of knowledge: Normal Help. It is surprising that participants in this condition were ambivalent about whether the CEO really knows that local water quality will be improved. After all, stopping pollution from being dumped into the water is an extremely reliable way of improving water quality. It is possible that random variation resulted in an unusually skeptical lot being assigned to this condition. But even if we suppose that the population mean for weighted knowledge ascription in Normal Help would be significantly above the midpoint at, say, 4, that would still be significantly below

32. $F(1, 222) = 10.54, p = .001, \eta_p^2 = .045$.

33. $F(2, 222) = 4.90, p = .008, \eta_p^2 = .042$.

34. Normal Harm, $M = 6.34, SD = 6.8$; Normal Help, $M = 0.03, SD = 3.82$; $t(63.17) = 3.45$, $MD = 6.31, p = 0.001$.

35. False Harm, $M = -2.68, SD = 8.64$; False Help, $M = -6.34, SD = 5.81$; $t(68.65) = 2.18$, $MD = 3.67, p = .033$.

36. Gettier Harm, $M = -6.67, SD = 6.04$; Gettier Help, $M = -5.88, SD = 6.35$; $t(79) = -0.57, p = .571$, n.s.

37. $t(37) = 5.733, p < .001$, test value = 0.

38. $t(33) = 0.02, p = .984$, n.s., test value = 0.

39. One-sample t-tests, all $ps \leq .057$, test value = 0.

40. Harm: 84%, Help: 50%, Fisher's exact test, $p = .002$.

41. Harm: 33%, Help: 11%, Fisher's exact test, $p = .051$.

42. Harm: 15%, Help: 14%, Fisher's exact test, $p = 1$.

43. Binomial test, $p < .001$.

44. Binomial test, $p = 1$.

45. Binomial tests, all $ps \leq .038$.

Figure 3: Experiment 5: *Top panel*: Percentage of participants ascribing knowledge in response to the dichotomous question. *Bottom panel*: Mean weighted knowledge ascriptions (scale ran -10 to +10).

what was observed in Normal Harm.⁴⁶ So, ultimately, I do not think that the unexpectedly low number in Normal Help should make us especially suspicious of the overall result.

The results from False Harm are also a bit surprising, with 33 % saying that the CEO really knows something that is false. Previous studies have shown that roughly 10–15 % of participants are typically willing to ascribe false knowledge when the belief is justified (Starmans and Friedman 2012; Buckwalter ms; Turri 2013), but the percentage observed here is significantly greater than even 15%.⁴⁷ By contrast, the percentage ascribing false knowledge in False Help is very similar to what has been observed previously in the literature.

Despite these two unexpected observations, there are no intolerable “red flags” in these results that should definitely make us reject the suggestion that participants are competently and literally applying their concept of knowledge. The overall results are mainly consistent with what mainstream theorists of knowledge would predict (again, the one exception being Normal Help). Yet we still observe a significant effect of side-effect valence. A constrained effect of this sort is arguably eligible to inform substantive theorizing about knowledge or our concept thereof.

7. Conclusion

By this point, it is definitely beyond reasonable doubt that the side-effect effect extends to knowledge ascriptions. ESEE is real and interesting in its own right. But what can it teach us? Ideally it would teach something about our concept of knowledge, and ultimately about knowledge itself. In order for it to teach us either of those things, the effect must manifest itself in a way that is plausibly due to a competent and literal application of our concept of knowledge. And in order to do that, it must be *appropriately constrained*. Up until now, that has not appeared to be the case (Experiments 1–4). But Experiment 5 provides evidence of an appropriately constrained effect.

My focus has been whether ESEE is appropriately constrained. But an equally important question for future research is whether it is *appropriately unconstrained*. At least three challenges lie ahead in this direction. I will briefly describe them here.

First, if our concept of knowledge, or knowledge itself, really is sensitive to evaluative facts, then this sensitivity will probably not be limited to side-effect propositions. A *side-effect proposition* is a proposition about a course of action’s side effects. For example, in the CEO case above, the side-effect proposition is that local water quality will be harmed or helped. A *central-effect proposition* is

46. $t(37) = 2.12, p = .041$, one-tailed, test value = 4.

47. Binomial test, $p = .004$, one-tailed, test proportion = .15.

a proposition about the primary effects that an action is intended to achieve. In the CEO case, the central-effect proposition is that the new program will increase company profits. It is hard to believe that our conceptual competence in ascribing knowledge would be sensitive to the moral valence of side effects but not to the moral valence of central effects. Similarly, if knowledge itself does depend on non-epistemically evaluative facts in surprising ways, it seems unlikely that this would be restricted to knowledge of side effects. Indeed, knowledge of many sorts of propositions, not necessarily connected to the consequences of the agent's actions, will probably likewise depend on evaluative facts.

Second, if our concept of knowledge, or knowledge itself, really is sensitive to evaluative facts, then this sensitivity will probably not be limited to knowledge based on testimony, or "second-hand knowledge." We should expect to see the effect in cases of "first-hand knowledge" too, such as beliefs based on perception, introspection, inference or memory. Perhaps this second challenge could be resisted on the grounds that different doxastic sources have very different profiles in folk epistemology,⁴⁸ in which case other factors might prevent a similar pattern for other sources.

Third, if our concept of knowledge, or knowledge itself, really is sensitive to evaluative facts, then this sensitivity will not be limited to a specific genre of cases. For example, it will not be limited to CEO cases. We should observe an appropriately constrained effect when using a wide range of cover stories. Prior work on ESEE has observed it using multiple cover stories, so I expect that this challenge will be met.

Those are three further tests for understanding ESEE's theoretical significance. If we do not observe an appropriately unconstrained effect, then ESEE might not manifest conceptual competence. By contrast, if we do observe a similar effect for central-effect propositions and first-hand knowledge and a wide range of cover stories, then three things follow. First, the effect will need a new name, since it would have nothing special to do with side effects. (Perhaps *the evaluative effect* will suffice.) Second, and more importantly, a conceptual-competence explanation of the effect becomes much more attractive. Third, consequently, substantive theorizing about the nature of knowledge will arguably have to take the effect into account.

Suppose those challenges are met and we are convinced that ESEE reveals something important about knowledge and our concept thereof. Then we will face an important set of further questions, including *why* knowledge is sensitive to these evaluative considerations, *how* this sensitivity is reflected in the psychology of knowledge attributions, and whether, upon reflection, our concept of knowl-

---

48. See Friedman and Turri (under review); Turri (forthcoming); and Turri (under review).

edge *should* be sensitive to such considerations. Why would knowledge work this way? How is that sensitivity reflected in the way people think about and ascribe knowledge? And once we are made aware that our ordinary practice of ascribing knowledge is sensitive in this way, is this something we should endorse, or should we instead change our concept so that it *is not* thus sensitive, or adopt a related concept that lacks such sensitivity? Some important progress has already been made along these lines by experimental philosophers and naturalists in epistemology (e.g. Schaffer 2008; Knobe 2010; Schaffer and Knobe 2012; Craig 1990; Kornblith 2002; Hawthorne 2004). But this research program would benefit tremendously from greater involvement from other areas of cognitive science besides epistemology and experimental philosophy, including cognitive anthropology, cognitive ethology, and cognitive, evolutionary and developmental psychology.

Appendix

The vignette for Experiment 4:

The vice president of a manufacturing company went to the CEO and said, “We are thinking of starting a new program. It will increase our profits, but there is a very, very small chance that it will harm local water quality, since it is just possible that it will require us to dump polluted water into the river behind our manufacturing plant. Nevertheless, that almost certainly won’t happen.” The CEO answered, “I agree that this plan will make us lots of profit. But I reject your claim that it is even possible that it will harm local water quality. Anyway, I don’t care at all about harming local water quality. I just want to make as much profit as I can.” Then the CEO made his decision: “Let’s start the new program and make some serious profits.” ¶ The company then started the new program, and sure enough, over the next few months, they made lots of profits. But the vice president was wrong about something: there was never a real chance that it would harm local water quality [But, against all odds, the program did harm local water quality].

Acknowledgments

For helpful conversation and feedback, I thank Mark Alfano, James Beebe, Peter Blouw, Wesley Buckwalter, Ori Friedman, Joshua Knobe, David Rose, Angelo Turri, and anonymous referees for *Ergo*. This research was supported by the Social

Sciences and Humanities Research Council of Canada and an Ontario Early Researcher Award.

References

- Adams, Fred, and Annie Steadman (2004). Intentional Action in Ordinary Language: Core Concept or Pragmatic Understanding? *Analysis*, 64(2), 173–181. <http://dx.doi.org/10.1093/analys/64.2.173>
- Alexander, Joshua, and Jonathan Weinberg (2007). Analytic Epistemology and Experimental Philosophy. *Philosophy compass*, 2(1), 56–80. <http://dx.doi.org/10.1111/j.1747-9991.2006.00048.x>
- Alicke, Mark D. (2008). Blaming Badly. *Journal of Cognition and Culture*, 8, 179–186. <http://dx.doi.org/10.1163/156770908X289279>
- Alicke, Mark D., and David Rose (2010). Culpable Control or Moral Concepts? *Behavioral and Brain Sciences*, 33(4), 330–331. <http://dx.doi.org/10.1017/S0140525X10001664>
- Austin, J. L. (1956–57). A Plea for Excuses. *Proceedings of the Aristotelian Society*, new series, 57, 1–30.
- Bach, Kent (2008). Applying Pragmatics to Epistemology. *Philosophical Issues*, 18, 68–88. <http://dx.doi.org/10.1111/j.1533-6077.2008.00138.x>
- Beebe, James, and Wesley Buckwalter (2010). The Epistemic Side-Effect Effect. *Mind and Language*, 25(4), 474–498. <http://dx.doi.org/10.1111/j.1468-0017.2010.01398.x>
- Beebe, James, and Mark Jensen (2012). Surprising Connections between Knowledge and Action: The Robustness of the Epistemic Side-Effect Effect. *Philosophical Psychology*, 25(5), 689–715. <http://dx.doi.org/10.1080/09515089.2011.622439>
- Beebe, James, and Joseph Shea (forthcoming). Gettierized Knobe Effects. *Episteme*.
- BonJour, Lawrence (2009). *Epistemology: Classical Problems and Contemporary Responses* (2nd ed.). Rowman and Littlefield.
- Buckwalter, Wesley (2012). Non-Traditional Factors in Judgments about Knowledge. *Philosophy Compass*, 7, 278–289. <http://dx.doi.org/10.1111/j.1747-9991.2011.00466.x>
- Buckwalter, Wesley (2013). Gettier Made ESEE. *Philosophical Psychology*. <http://dx.doi.org/10.1080/09515089.2012.730965>
- Buckwalter, Wesley (n.d.). Factive Verbs and Protagonist Projection. Unpublished manuscript.
- Buckwalter, Wesley, David Rose, and John Turri (2013). Belief Through Thick and Thin. *Noûs*. <http://dx.doi.org/10.1111/nous.12048>
- Burge, Tyler (1979). Individualism and the Mental. *Midwest Studies in Philosophy*, 4, 73–121. <http://dx.doi.org/10.1111/j.1475-4975.1979.tb00374.x>
- Burge, Tyler (1986). Individualism and Psychology. *Philosophical Review*, 95, 3–45. <http://dx.doi.org/10.2307/2185131>
- Cohen, Stewart (1999). Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13, 57–89.
- Craig, Edward (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford University Press.
- Cullen, Simon (2010). Survey-Driven Romanticism. *Review of Philosophy and Psychology*, 1, 275–296. <http://dx.doi.org/10.1007/s13164-009-0016-1>
- Davidson, Donald (2001). A Coherence Theory of Truth and Knowledge. In Donald Da-

- vidson (Ed.), *Subjective, Intersubjective, Objective* (137–157). Oxford University Press. (Reprinted from *Kant oder Hegel?*, 423–438, by Dieter Henrich, Ed., 1983, Klett-Cotta).
- Davis, Wayne (2007). Knowledge Claims and Context: Loose Use. *Philosophical Studies*, 132, 395–438. <http://dx.doi.org/10.1007/s11098-006-9035-2>
- DeRose, Keith (2009). *The Case for Contextualism: Knowledge, Skepticism, and Context* (Vol. 1). Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199564460.001.0001>
- Fantl, J., & McGrath, M. (2002). Evidence, Pragmatics, and Justification. *The Philosophical Review*, 111(1), 67–94.
- Fantl, Jeremy, and Matthew McGrath (2009). *Knowledge in an Uncertain World*. Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199550623.001.0001>
- Feldman, Richard (2003). *Epistemology*. Prentice Hall.
- Feltz, Adam, and Chris Zarpentine (2010). Do You Know More When It Matters Less? *Philosophical Psychology*, 23, 683–706. <http://dx.doi.org/10.1080/09515089.2010.514572>
- Friedman, Ori, and John Turri (under review). Is Probabilistic Evidence a Source of Knowledge? Manuscript submitted for publication.
- Fumerton, Richard (2006). *Epistemology*. Blackwell.
- Gettier, Edmund (1963). Is Justified True Belief Knowledge? *Analysis*, 23(6), 121–123. <http://dx.doi.org/10.1093/analys/23.6.121>
- Guglielmo, Steve, and Bertram Malle (2010). Can Unintended Side Effects be Intentional? Resolving a Controversy over Intentionality and Morality. *Personality and Social Psychology Bulletin*, 36(12), 1635–1647. <http://dx.doi.org/10.1177/0146167210386733>
- Hawthorne, John (2004). *Knowledge and Lotteries*. Oxford University Press.
- Hazlett, Allan (2010). The Myth of Factive Verbs. *Philosophy and Phenomenological Research*, 80, 497–522. <http://dx.doi.org/10.1111/j.1933-1592.2010.00338.x>
- Hazlett, Allan (forthcoming). Factive Presupposition and the Truth Condition on Knowledge. *Acta Analytica*.
- Holton, Richard (1997). Some Telling Examples: A Reply to Tsohatzidis. *Journal of Pragmatics*, 28, 624–628. [http://dx.doi.org/10.1016/S0378-2166\(96\)00081-1](http://dx.doi.org/10.1016/S0378-2166(96)00081-1)
- Lewis, David (1996). Elusive Knowledge. *Australasian Journal of Philosophy*, 74, 549–567. <http://dx.doi.org/10.1080/00048409612347521>
- Knobe, Joshua (2003a). Intentional Action and Side Effects in Ordinary Language. *Analysis*, 63, 190–193. <http://dx.doi.org/10.1093/analys/63.3.190>
- Knobe, Joshua (2003b). Intentional Action in Folk Psychology: An Experimental Investigation. *Philosophical Psychology*, 16, 309–324. <http://dx.doi.org/10.1080/09515080307771>
- Knobe, Joshua (2004). Intention, Intentional Action and Moral Considerations. *Analysis*, 64, 181–187. <http://dx.doi.org/10.1093/analys/64.2.181>
- Knobe, Joshua (2010). Person as Scientist, Person as Moralist. *Behavioral and Brain Sciences*, 33, 315–329. <http://dx.doi.org/10.1017/S0140525X10000907>
- Knobe, Joshua, and Arudra Burra (2006). Experimental Philosophy and Folk Concepts: Methodological Considerations. *Journal of Cognition and Culture*, 6(1–2), 331–342. <http://dx.doi.org/10.1163/156853706776931402>
- Kornblith, Hilary (2002). *Knowledge and Its Place in Nature*. Oxford University Press. <http://dx.doi.org/10.1093/0199246319.001.0001>
- Kripke, Saul (1972). *Naming and Necessity*. Harvard University Press.

- Lehrer, Keith (2000). *Theory of Knowledge* (2nd ed.). Westview.
- Malle, Bertram (2006). Intentionality, Morality, and their Relationship in Human Judgment. *Journal of Cognition and Culture*, 6(1–2), 87–112. <http://dx.doi.org/10.1163/156853706776931358>
- Matilal, Bimal (1986). *Perception: An Essay on Classical Indian Theories of Knowledge*. Oxford University Press.
- Murray, Dylan, Justin Sytsma, and Jonathan Livengood (2012). God Knows (but Does God Believe?). *Philosophical Studies*. <http://dx.doi.org/10.1007/s11098-012-0022-5>
- Nadelhoffer, Thomas (2006). Bad Acts, Blameworthy Agents, and Intentional Actions: Some Problems for Jury Impartiality. *Philosophical Explorations*, 9, 203–220. <http://dx.doi.org/10.1080/13869790600641905>
- Nagel, Jennifer (2008). Knowledge Ascriptions and the Psychological Consequences of Changing Stakes. *Australasian Journal of Philosophy*, 86, 279–294. <http://dx.doi.org/10.1080/00048400801886397>
- Pettit, Dean, and Joshua Knobe (2009). The Pervasive Impact of Moral Judgment. *Mind & Language*, 24, 586–604. <http://dx.doi.org/10.1111/j.1468-0017.2009.01375.x>
- Phillips, Stephen (2011). Epistemology and Classical Indian Philosophy. *Stanford Encyclopedia of Philosophy*, Spring 2011 edition. Ed. Edward Zalta. Accessed 15 Sept. 2012. Retrieved from <http://plato.stanford.edu/archives/spr2011/entries/epistemology-india/>
- Putnam, Hilary (1975). The Meaning of ‘Meaning’. *Philosophical Papers: Vol. 2. Mind, Language, and Reality* (215–271). Cambridge University Press.
- Radford, C. (1966). Knowledge – By Examples. *Analysis*, 27(1), 1–11.
- Rose, David, and Jonathan Schaffer (2013). Knowledge Entails Dispositional Belief. *Philosophical Studies*. <http://dx.doi.org/10.1007/s11098-012-0052-z>
- Rysiew, Patrick (2001). The Context-Sensitivity of Knowledge Attributions. *Noûs*, 35(4), 477–514. <http://dx.doi.org/10.1111/0029-4624.00349>
- Santayana, George (1923). *Skepticism and Animal Faith*. Dover, 1955.
- Sartwell, Crispin (1991). Knowledge is Merely True Belief. *American Philosophical Quarterly*, 28(2), 157–165.
- Schaffer, Jonathan (2008). Knowledge in the Image of Assertion. *Philosophical Issues*, 18, 1–19. <http://dx.doi.org/10.1111/j.1533-6077.2008.00134.x>
- Schaffer, Jonathan, and Joshua Knobe (2012). Contrastive Knowledge Surveyed, *Noûs*, 46, 675–708. <http://dx.doi.org/10.1111/j.1468-0068.2010.00795.x>
- Schwitzgebel, Eric, and Blake Myers-Schulz (forthcoming). Knowing that P without Believing that P. *Noûs*.
- Sosa, Ernest (1991). *Knowledge in Perspective: Selected Essays in Epistemology*. Cambridge University Press. <http://dx.doi.org/10.1017/CBO9780511625299>
- Sripada, Chandra, and Jason Stanley. (2012). Empirical Tests of Interest-Relative Invariantism. *Episteme*, 9(1), 3–26. <http://dx.doi.org/10.1017/epi.2011.2>
- Stanley, Jason (2005). *Knowledge and Practical Interests*. Oxford University Press. <http://dx.doi.org/10.1093/0199288038.001.0001>
- Starmans, Christina, and Ori Friedman (2012). The Folk Conception of Knowledge. *Cognition*, 124, 272–283. <http://dx.doi.org/10.1016/j.cognition.2012.05.017>
- Steup, Matthias (1996). *An Introduction to Contemporary Epistemology*. Prentice Hall.
- Swain, Stacey, Joshua Alexander, and Jonathan Weinberg (2008). The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp. *Philosophy and*

- Phenomenological Research*, 76(1), 138–155. <http://dx.doi.org/10.1111/j.1933-1592.2007.00118.x>
- Turri, John (2011). Manifest Failure: The Gettier Problem Solved. *Philosophers' Imprint*, 11(8), 1–11.
- Turri, John (2012a). Is Knowledge Justified True Belief? *Synthese*, 184(3), 247–259. <http://dx.doi.org/10.1007/s11229-010-9773-8>
- Turri, John (2012b). In Gettier's Wake. *Epistemology: The Key Thinkers* (214–229). Ed. Stephen Hetherington. Continuum.
- Turri, John (2013). A Conspicuous Art: Putting Gettier to the Test. *Philosophers' Imprint*, 13(10), 1–16.
- Turri, John (2014). Linguistic Intuitions in Context: A Defense of Pure Nonskeptical Invariantism. *Intuitions*. Ed. Anthony Booth and Darrell Rowbottom. Oxford University Press.
- Turri, John (forthcoming). Skeptical Appeal: The Source-Content Bias and Its Application to Skepticism. *Cognitive Science*.
- Turri, John. (under review). An Open-and-Shut Case: Epistemic Closure in the Manifest Image. Manuscript submitted for publication.
- Ulatowski, J. (2012). Act Individuation: An Experimental Approach. *Review of Philosophy and Psychology*, 3, 249–262. <http://dx.doi.org/10.1007/s13164-012-0096-1>
- Weinberg, Jonathan, Shaun Nichols, and Stephen Stich (2001). Normativity and Epistemic Intuitions. *Philosophical Topics*, 29, 429–460. <http://dx.doi.org/10.5840/philtopics2001291/217>
- Young, Liane, Fiery Cushman, Ralph Adolphs, Daniel Tranel, and Marc Hauser (2006). Does Emotion Mediate the Effect of an Action's Moral Status on its Intentional Status? Neuropsychological Evidence. *Journal of Cognition and Culture*, 6, 265–278. <http://dx.doi.org/10.1163/156853706776931312>
- Zagzebski, Linda (1994). "The Inescapability of Gettier Problems." *The Philosophical Quarterly*, 44(174), 65–73. <http://dx.doi.org/10.2307/2220147>
- Zagzebski, Linda (1996). *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press. <http://dx.doi.org/10.1017/CBO9781139174763>