Skeptical appeal: the source-content bias[*]

John Turri
john.turri@gmail.com

Abstract: Radical skepticism is the view that we know nothing, or at least next to nothing. Nearly no one actually believes that skepticism is true. Yet it has remained a serious topic of discussion for millennia and it looms large in popular culture. What explains its persistent and widespread appeal? How does the skeptic get us to doubt what we ordinarily take ourselves to know? I present evidence from two experiments that classic skeptical arguments gain potency from an interaction between two factors. First, people evaluate inferential belief more harshly than perceptual belief. Second, people evaluate inferential belief more harshly when its content is negative (i.e. that something is *not* the case) than when it's positive (i.e. that something *is* the case). It just so happens that potent skeptical arguments tend to focus our attention on *negative inferential* beliefs, and we are especially prone to doubt that such beliefs count as knowledge. That is, our cognitive evaluations are biased against this specific combination of source and content. The skeptic sows seeds of doubt by exploiting this feature of our psychology.

1. Introduction

We face uncertainty all the time. Doubt and ignorance are unavoidable facts of life. But they are neither permanent nor irreparable. As much as doubt and ignorance dot the landscape of our existence, so too does knowledge. Inquiry remedies ignorance and diminishes uncertainty: we pass, step by step, from ignorance to knowledge, from uncertainty to understanding. In this respect, modern science is a macrocosm of the cognitive path we each take through life. Our cognitive

[*] This is the penultimate version of a paper to appear in *Cognitive Science*. Please cite the final, published version if possible.

powers are fallible and limited, but no less impressive for all that.

Skeptics notoriously disagree. Radical skepticism is the view that we know nothing, or at least next to nothing. Some forms of skepticism allow that we know facts about how things currently appear to us (e.g. "I seem to be in pain") or obvious basic conceptual truths (e.g. "2 > 1"). But, in general, to be skeptical about a domain is to deny that we know that specific propositions in that domain are true. Many researchers assume that knowledge is a mental state (Premack & Woodruff, 1978; Pillow, 1989; Povinelli, Nelson & Boysen, 1990; Perner, Frith, Leslie & Leekam, 1989; Povinelli & deBlois, 1992; Flavell, 2000; Williamson, 2000; Hare, Call & Tomasello, 2001; Sprung, 2010). Accordingly, we can frame skepticism in terms of mental state ascription: to be skeptical about a domain is to deny all knowledge ascriptions regarding that domain. People might have beliefs or convictions or hopes or intentions about the domain, but they don't have *knowledge*.

Skepticism lives a life beyond Philosophy 101. Skepticism readily captures the public imagination and is a staple of popular culture. Gaming boards feed on it. For instance, the motto of the New York State Lottery is, "Hey, you never know" (http://nylottery.ny.gov/wps/portal). Hollywood grows rich on it. Witness the many successful films that put skeptical doubts front and center. Foremost among these are *The Matrix* (do you know that you're not in The Matrix?), *Vanilla Sky* (do you know that you're life is not just a dream?), *The Truman Show* (do you know that your life isn't just one big charade?), *Bladerunner* (do you know that this individual isn't a replicant?), *Dark City* and *Memento* (do you know that you're past is what it seems to be?). If nothing else, skillful trafficking in skeptical doubt is a lucrative business. Not only does skepticism live a life beyond Philosophy 101, but apparently the skilled skeptic can live lavishly.

Nevertheless, just as no battle plan survives contact with the enemy, radical skepticism doesn't survive contact with daily life. The skeptic tells us that we

don't know that we're not merely dreaming, or that other people have minds, or that the universe is older than five minutes, or that our past is what it seems to be, or that there is an objective world outside our minds, or that snow is white, or that the next time we eat bread it will nourish us, etc. (Sextus Empiricus; Descartes, 1641; Bayle, 1702; Berkeley, 1710; Hume, 1748; Russell, 1921; Unger, 1975). We find these assertions incredible and are prone to reject them out of hand as ridiculous and inconsistent with common sense (Reid, 1764). Sane people do not and cannot seriously doubt such things in practice. But even granting that this is the wise response, residual dissatisfaction remains. Unless we're able to identify a flaw in the skeptic's performance — some error or trick that takes us in — thoughtful people will continue to find the situation somewhat troubling (Stroud, 1984; DeRose, 1995). True, we carry on with our lives in much the same way (Hume, 1739: Bk 1.4.7); but it is with a slightly heavier heart. No matter if we can't figure out *how* the peddler swindles people with his shell game, the wise person just walks right by and keeps her money in her pocket. But the *curious* wise person still feels the mild sting of ignorance and suffers more than a modicum of intellectual anxiety.

All this raises an interesting psychological question. If nearly no one accepts skepticism, then why do skeptical arguments perennially fascinate us? Why do we, individually and collectively, over thousands of years, keep revisiting skepticism?

This paper begins the task of explaining what in our psychology renders us vulnerable to these skeptical doubts. This entails a subtle shift in focus from the recent flurry of valuable experimental work on commonsense cognitive evaluations (Swain, Alexander & Weinberg, 2008; Wright, 2010; May et al., 2010; Beebe & Buckwalter, 2010; Feltz & Zarpentine, 2010; Starmans & Friedman, 2012; Schaffer & Knobe, 2012; Beebe & Jensen, 2012; Sripada & Stanley, 2012; Turri, 2013; Nagel, San Juan & Mar, 2013; Myers-Shulz & Schwitzgebel, 2013;

Murray, Sytsma & Livengood, in press; Buckwalter & Schaffer, in press; Buckwalter, Rose & Turri, in press), otherwise known as "folk epistemology" (Montgomery 1992; Bartsch 2002; Hardy-vallée & Dubreuil 2010; Mercier 2010; Spicer 2010; Turri & Friedman 2014). While knowledge attribution has been the focus of considerable attention recently, its conceptual flip-side has not: knowledge *denial*. Knowledge denial is the skeptic's stock in trade. So our guiding question might be put this way: how does the skeptic get us interested in purchasing his wares?

1.1. The classical argument for skepticism and two examples

We will focus on the most celebrated and influential skeptical strategy (Descartes, 1641; Moore, 1959; Dretske, 1970; Unger, 1975; Stine, 1976; Goldman, 1976; Lewis, 1979; Nozick, 1981;Stroud, 1984; Cohen, 1988; DeRose, 1995; Sosa, 1999; Pritchard, 2002; Cohen, 2005; for an overview see Turri 2014, esp. §§1-4, 19-23, 47-49). The strategy is to proceed in three stages. *First* identify some claim that we would ordinarily take ourselves to know (call it 'O' for '*o*rdinary claim'). *Next* identify some tricky claim (call it 'T' for '*t*ricky claim'), the *denial* of which (~T) obviously follows from O. It's important that we have the intuition that we ordinarily could *not* know ~T. *Then* reason as follows.

1. If we know that O, then we know that ~T.
2. But we don't know that ~T.
3. So we don't know that O.

This is sometimes called "the classical argument for skepticism" and characterized as "exceedingly compelling" (Unger, 1975: 8–9). The argument is logically valid and appeals to premises that we find plausible. Yet its conclusion is incredible and contradicts something we normally take for granted. Some consider this to be a genuine *paradox* (Cohen, 1988). For decades resolving this paradox has

been a central concern in the theory of knowledge.

This classic argumentative strategy can be used to challenge many, if not most, of the things we ordinarily take ourselves to know (Vogel, 1990; Hawthorne, 2004). We'll briefly consider two examples familiar from the literature on skepticism, which will then inform the experimental studies reported below.

First, we ordinarily take ourselves to know where various objects important to us are. For instance, when Maxwell is done at work for the day, he knows that his car is parked in the parking garage (O). However, it's not uncommon for cars to be stolen. Does Maxwell know that his car has *not* been stolen (~T)? It can easily seem that Maxwell does not know this (Vogel, 1990). But, of course, Maxwell does know that if his car is parked in the garage, then it hasn't been stolen (if O, then ~T). So, the skeptic invites us to conclude, Maxwell doesn't know where his car is after all.

Second, we ordinarily take ourselves to know what natural kinds things belong to. For instance, if Zoe goes to the zoo and sees a large, menacing, spotted feline in the exhibit labelled "Jaguar," she knows that the animal is a jaguar (O). However, it's not unheard of for zoos to make mistakes or temporary substitutions in a pinch to please their patrons. Does Zoe know that the animal is *not* a leopard that looks just like a jaguar in a jaguar display (~T)? It can seem that Zoe does not know this (Dretske, 1970: 1016; Unger, 1975: 8). But, of course, Zoe knows that if the animal is a jaguar, then it's not a leopard of any sort. More pointedly, she knows that if it's a jaguar, then it's not a leopard that looks just like a jaguar in a jaguar display (if O, then ~T). But Zoe does not know that the animal is *not* a leopard that looks just like a jaguar in a jaguar display. So, the skeptic invites us to conclude, Zoe doesn't know that it's a jaguar after all.

We could multiply examples, but the two we've covered are enough to illustrate how the skeptic's strategy works. The experiments reported below use vi-

gnettes based on these two examples familiar from the literature on skepticism.

1.2. A theory in two parts

Here is one theory about why the classical skeptical argument can seem so potent. The theory has two parts. On the one hand, people view perception as a more likely source of knowledge than inference. That is, other things being equal, people more readily classify a belief as knowledge when it's based on observation than when it's based on inference from background knowledge. On the other hand, people view *positive* inferential belief as a more likely candidate for knowledge than *negative* inferential belief. That is, other things being equal, an inferential belief that something *is the case* is more likely to be classified as knowledge than an inferential belief that something *is not the case*.

If true, these two points can explain the classical skeptical argument's force. First, the argument gets us to focus on what seems like an inferential belief. We're left to think that *observation* is of no further help, so the agent in question is left to *infer* the relevant proposition from prior experience and background knowledge. Second, the argument gets us to focus on whether something is *not* the case, so the inferential belief in question has *negative content*. On this theory, it's possible that the appeal of skepticism has nothing special to do with either the form of argument that the skeptic employs or the cleverness of the skeptic's tricky alternative scenarios. Rather, it could just be that skepticism's appeal is due to the fact that we're heavily biased against classifying negative inferential beliefs as knowledge.

Both experiments below were designed to help evaluate these possibilities. Experiment 1 examines the effect of two factors on knowledge denial in contexts where skeptical possibilities are explicitly raised and considered. The two factors are a belief's source, perceptual versus inferential, and a belief's content, positive

versus negative. If my hypothesis is correct, then we should see high rates of knowledge denial precisely in cases of negative inferential belief. This is in fact what we observe. Experiment 2 focuses in on inferential belief specifically and again examines the effect of negative versus positive content on people's willingness to deny knowledge. Again we see that negative inferential belief is evaluated more harshly than positive inferential belief, even when each is based on a long and unbroken track record of evidence.

Experiment 2 also accomplishes something further. The skeptic focuses only on the ordinary claim (O) and the denial of the tricky alternative scenario (~T); he draws skeptical conclusions based this comparison and invites us to do the same. It might seem that the difficulty here stems from something about denying *this particular alternative*. But the difficulty might instead stem from *denying* this particular alternative. To test between these possibilities, Experiment 2 examines rates of knowledge attribution for O and ~T alongside those for the denial of O (i.e. ~O). The results suggest that skeptic induces doubt not because his alternative is especially difficult to rule out, but rather because the focus shifts from affirmative to negative content.

My hypothesis about the source of skepticism's allure can be motivated in light of two substantial bodies of psychological research. On the one hand, researchers have found that people exhibit an "anti-inference bias" and are more comfortable relying on direct perceptual evidence, even when the perceptual evidence is acknowledged to be less reliable than the inference (Zamir, Rito & Teichman, in press). Young children are apt to attribute knowledge to subjects with perceptual access to a fact, (Pillow, 1989), but they are often reluctant to attribute knowledge based on simple valid inferences (Sodian & Wimmer 1987). Evidence from comparative psychology suggests that nonhuman animals, including chimpanzees (Povinelli & Eddy, 1996; Call & Tomasello, 2008), bottlenose dolphins (Reiss, 2001), domestic dogs (Horowitz, 2008), and several bird species (Dally,

Emery & Clayton, 2005; Keefner, 2013) attribute mental states to conspecifics, and that these attributions are influenced by what the conspecific has perceptual access to. On the other hand, decades of behavioral research show that it requires more effort to represent and perform computations over negations than affirmations (see Conclusion for references). Moreover, increased effort can be expected to diminish confidence in the truth of a claim (Reber & Schwarz, 1999; Oppenheimer, 2008), and since knowledge requires truth, this could in turn be expected to increase rates of knowledge denial.

In addition to the fact that my hypothesis can be grounded in prior psychological findings, it is novel in two respects. First, it marshals those findings in order to diagnose the psychological source of a specific, perennial philosophical conundrum.[1] Second, it posits an *interaction effect* of a belief's *source* and *content* on judgments about knowledge. While prior research has noted the effects of belief source and negative content, respectively, the interaction of these two factors has not been documented. Prior influential work in cognitive science has made progress in identifying psychological mechanisms and processes by studying moral cognition through the lens of longstanding philosophical disputes in ethics, especially those motivated by poignant thought experiments involving moral dilemmas (e.g. Greene et al. 2001; Cushman & Greene 2012). If my hypothesis is correct, then the same could be true of studying cognitive evaluations through the lens of classical skeptical challenges.

[1] Thanks to an anonymous reviewer for suggesting language along these lines.

2. Experiment 1

2.1. Method

2.1.1. Participants

Participants (N = 607, 204 female, aged 18–69, mean age = 29.68 years) were recruited and tested using an online platform (Amazon Mechanical Turk + Qualtrics) and compensated $0.30 for approximately 2 minutes of their time. Participation was restricted to those residing in the United States and 97 percent reported English as a native language. Participants were not allowed to re-take the survey; repeat participation was prevented by screening AMT Worker IDs. They filled out a brief demographic survey after testing. I excluded data from 32 participants who failed comprehension questions. Including data from these participants didn't affect the pattern of results reported below.

2.1.2. Materials and procedure

Participants were randomly assigned to one of eight conditions in a 2 (Story: Car, Specimen) × 2 (Content: Negative, Positive) × 2 (Source: Inference, Perception) between-subjects design. Participants in each condition read a different story. Two different storylines were used to ensure that the results weren't due to superficial features of a particular case. The stories were based on widely discussed examples in the literature on skepticism (see section 1.1). For each storyline, the source and content of the protagonist's belief was varied similarly.

The basic storyline for Car featured an office worker, Maxwell, who parked his car in the parking garage when he arrived at work a while ago. Now it's lunchtime and Maxwell and his assistant are looking for a document. Maxwell suspects that he left the document in his car. Maxwell's assistant asks him where his car is parked and, in the process, points out, "It's not unheard of for cars to be

stolen." The Content factor varies which aspect of the assistant's speech Maxwell focuses on: *where the car is parked* versus *whether the car has been stolen*. Participants in the Negative conditions were asked whether Maxwell knows that his car was *not* stolen. Participants in the Positive conditions were asked whether Maxwell knows that his car *is* parked in the spot he parked in this morning. The Source factor varies what source Maxwell bases his answer on. In the Inference conditions, Maxwell *thinks for a moment* before answering, unaided by observation. In the Perception condition, Maxwell *looks out the window* before answering. Here is the story (manipulations bracketed and separated by a slash):

> When Mr. 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 for a moment / looks out the window] and then responds, "[No, my car hasn't been stolen / Yes, my car is in C8]."

After reading the story, participants answered two comprehension questions and a test question about what Maxwell knows (options in parentheses)

1. Maxwell parked in _____. (C8/D8)
2. Maxwell is in the _____. (archives room/parking garage/cafeteria)
3. Maxwell _____ that his car [is parked in C8 / has not been stolen]. (knows/only believes)

The basic storyline for Specimen featured an avid zoo patron, Michelle, whose favorite exhibit is "The Big Cat Exhibit." Over thousands of observations, the animal in the exhibit has always been a jaguar. The Source factor varies

whether Michelle judges the animal based on perception (Perception) or based on an inference unaided by observation (Inference). The Content factor varies whether Michelle thinks that the animal *is* a jaguar (Positive), or that the animal is *not* a leopard (Negative). Participants in the Positive condition were asked whether Michelle knows that the animal is a jaguar; in the Negative condition they were asked whether Michelle knows that it's not a leopard. Here is the story:

> Michelle has visited the city zoo every day for the past ten years. Her favorite exhibit is The Big Cat Exhibit. Over thousands of observations, the animal in this exhibit has always been a jaguar. Today [Michelle must stay home and can't visit the zoo because she sprained her ankle. While relaxing on the couch / when Michelle left home to visit the zoo, she almost sprained her ankle. While looking at the animal in the exhibit], Michelle thinks, "The animal in the Big Cat Exhibit today is [a jaguar / not a leopard]." And she is right: it is [a jaguar / not a leopard].

After reading the story, participants were asked two comprehension questions and a test question:

1. The animal _____ a jaguar. (is/is not)
2. Michelle is _____. (at home/at the zoo)
3. Michelle _____ that the animal is [a jaguar / not a leopard]. (knows/only believes)

Participants who failed a comprehension question were excluded from the analysis. Questions were always asked in the same order and response options were always rotated randomly. The story remained at the top of the page throughout the survey.

2.2. Results

An initial omnibus regression was used to determine whether knowledge denial

was predicted by storyline or participant gender. In line with previous research on knowledge ascriptions (Wright, 2008; Turri, 2013; Nagel et al., 2013; Friedman and Turri, under review), participant gender didn't predict knowledge denial. There was no main effect of participant gender (Wald = 0.012, df = 1, p = .911, n.s., all tests two-tailed unless otherwise noted) and no interaction with Content (Wald = 2.11, df = 1, p = .146), Source (Wald = 0.03, df = 1, p = .863), or Story (Wald = 0.534, df = 1, p = .465). Similarly, logistic regression showed that there was no main effect of Story (Wald = 1.161, df = 1, p = .281, n.s.) and no interaction with Content (Wald = 1.03, df = 1, p = .310, n.s.) or Source (Wald = 0.021, df = 1, p = .885). Because gender and story had no effect, the analyses that follow collapse across these factors.

Binary logistic regression was performed to assess the impact of Source and Content on the likelihood that participants would deny knowledge (Table 1). The model contained Source, Content, and the interaction of Source and Content as predictor variables and knowledge denial as the outcome variable. (Preliminary analysis was conducted to ensure no violation of multicollinearity assumptions.) In light of the hypothesis proposed in section 1.2, I expected two things. First, there would be a main effect of Source, with inferential belief predicting higher rates of knowledge denial. Second, there would be a significant interaction between Source and Content, with negative content more strongly predicting higher rates of knowledge denial for inferential belief.

Both predictions were true (Table 1). The full model was statistically significant, $\chi^2(3, N = 607) = 93.82$, $p < .001$. It explained between 14.3% (Cox and Snell R Square) and 19.4% (Nagelkerke R Square) of the variance in knowledge denial, and it correctly classified 70% of cases. Source of belief made a unique statistically significant contribution to the model. By changing the source of the protagonist's belief from perception to inference, the odds of denying knowledge increased by 280% (or a factor of 2.8). Content of belief did not make a unique sta-

tistically significant contribution to the model. Crucially, there was a significant interaction between the source and content of belief: by changing the content of the protagonist's belief from positive to negative, the odds of denying knowledge increased more for inferential belief than for perceptual belief (by a factor of 2.33).

Table 1. Experiment 1. Logistic regression predicting knowledge denial. (Reference class for Source: Perception. Reference class for Content: Positive.)

	B	S.E.	Wald	df	*p*	Odds Ratio	95% C.I. for Odds Ratio	
							Lower	Upper
Source	1.03	.26	16.13	1	<.001	2.8	1.70	4.64
Content	.349	.27	1.69	1	=.194	1.42	.84	2.40
Source*Content	.847	.36	5.43	1	=.02	2.33	1.14	4.76
Constant	-1.33	.20	44.77	1	<.001	.264		

As a further way of clarifying and illustrating the results, I report some pairwise comparisons among the cells. If the hypothesis proposed in section 1.2 is correct, then we would expect three things. We would expect knowledge denial to be (1) significantly higher in Negative Inference than in Positive Inference, (2) significantly higher in Positive inference than in Positive Perception and Negative Perception, but (3) no different between between Positive Perception and Negative Perception.

All three expectations are met (Figure 1). First, rate of knowledge denial was higher in Negative Inference (N = 145, 71%) than in Positive Inference (N = 155, 42.6%), Fisher's exact test, $p < .001$, Cramer's V = .287. Second, knowledge denial in Positive Inference was higher than in Positive Perception (N = 153, 20.9%), Fisher's exact test, $p < .001$, Cramer's V = .233, as well as Negative Perception (N = 155, 27.3%), Fisher's exact test, $p = .006$, Cramer's V = .161. Third, rate of knowledge denial didn't differ between Positive Perception and Negative

Perception, Fisher's exact test, $p = .230$, n.s.

Finally, binomial tests revealed that only in Negative Inference did participants deny knowledge at rates exceeding chance ($p < .001$). By contrast, knowledge denial was trending below chance in Positive Inference ($p < .077$) and fell far below chance in both Positive Perception and Negative Perception ($ps < .001$).

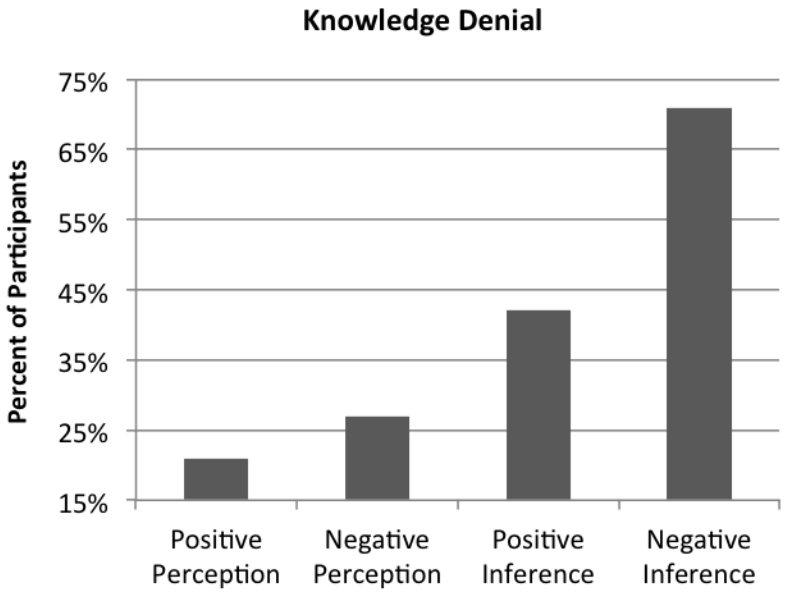

Fig. 1. Experiment 1. Denial of knowledge across conditions.

2.3. Discussion

Three main findings emerge from this study. First, participants were more likely to deny knowledge in cases of inference than in cases of perception. An inferential belief is much less likely than a perceptual belief is to be classified as knowledge. Second, a negative inferential belief is much less likely than a positive inferential belief is to be classified as knowledge. The cumulative result is that negative inferential belief suffers a double-deficit. The interaction effect in the regression model is particularly good evidence for this double-deficit. (For replications

of this interaction effect, see Turri under review.) Third, only when evaluating a negative inferential belief did participants deny knowledge at rates exceeding chance. In each of the other three conditions — when evaluating positive perceptual belief, negative perceptual belief, and positive inferential belief — participants *ascribed* knowledge at rates exceeding chance.

Importantly, we observed these patterns while using stories directly based on influential cases from the contemporary literature on skepticism. Thus the results suggest that skeptics instill doubt in us to the extent that they get us to focus on inferential belief with negative content. Skeptical arguments themselves could function simply to shift our focus in that way, rather than persuading us in virtue of their logic.

These results suggest that the skeptic's alternative scenarios — "your car might have been stolen," "the animal might be a leopard that looks just like a jaguar" — aren't especially challenging to rule out. That is, it's possible that the skeptic has simply alighted on a class of beliefs that we're antecedently inclined to evaluate especially harshly. A second study was designed to investigate this possibility further by examining whether ruling out the skeptic's alternatives is any more difficult than ruling out more routine claims.

3. *Experiment 2*

The skeptic singles out the *logical implications* of propositions we ordinarily take ourselves to know. If we ordinarily take ourselves to know O, and we know that O entails ~T, then we should also at least be in a position to know ~T. For example, if the animal is a jaguar, then it's not a leopard of any sort. So if you know that the animal is a jaguar, then you should also be in a position to know that it's not a leopard. By contrast, the skeptic doesn't consider whether knowing O is any easier than knowing *its negation* (~O). If the skeptic really has identified a special

problem, then it should be harder to know ~T than it is to know ~O (other things being equal). In other words, it's possible that the key difference might not be between O and T but rather between negative and affirmative contents.

The Specimen cover story is particularly well suited to investigate this. We ordinarily take ourselves to know that the animal in the display is a jaguar. And, obviously, if it's a jaguar, then it's not a leopard. But jaguars and leopards look very similar. Do we know that it's not a leopard? As we saw in Experiment 1, when the agent is able to *observe* the animal, people tend to say that she knows that it's not a leopard. By contrast, when she doesn't observe the animal but instead draws an inference, people deny that she knows that it's not a leopard. People denied knowledge even though her inference was based on a long and unbroken track record: over many years and thousands of observations, the animal in the display has always been a jaguar.

But suppose that instead of always being a jaguar, the animal in the display had *never* been a jaguar. Based on this long and unbroken track record, will people be any less harsh in evaluating an inferential belief that the animal is *not a jaguar*?

If the skeptic has identified a special problem, then it should be viewed as *harder* to know

> (~T) that the animal is not a leopard (based on a long and unbroken track record of its being a jaguar)

than it is to know

> (~O) that the animal is not a jaguar (based on a long and unbroken track record of its not being a jaguar).

By contrast, if my earlier hypothesis is correct, then people will view those two propositions as *equally difficult* to know. In other words, rates of knowledge denial for those two claims will be equally high because they are both inferential

and negative. Moreover, rates of knowledge denial for those two claims should be significantly higher than for this claim, which is inferential but positive:

> (O) that the animal is a jaguar (based on a long and unbroken track record of its being a jaguar).

The present study directly tests these suggestions.

3.1. Method

3.1.1. Participants

Three hundred and five new participants (87 female, aged 18-78, mean age = 28.16 years) were tested. They were recruited and compensated the same way as in Experiment 1. Ninety-eight percent reported English as a native language. I excluded data from 11 participants who failed comprehension questions. Including data from these participants didn't affect the pattern of results reported below.

3.1.2. Materials and procedure

Participants were randomly assigned to one of three conditions: Positive Jaguar, Negative Leopard, and Negative Jaguar. Participants in each condition read a different story. The stories for the Positive Jaguar and Negative Leopard conditions in this experiment are very similar to the Specimen stories for the Positive Inference and Negative Inference conditions in Experiment 1. In the story for Positive Jaguar, the animal has always been a jaguar and Michelle thinks that it's a jaguar today too. In the story for Negative Leopard, the animal has always been a jaguar and Michelle thinks that it's not a leopard today either. The story for the Negative Jaguar condition is similar except that this time the animal in the display has *never* been a jaguar and Michelle thinks that it's not a jaguar today either:

> Michelle has visited the city zoo every day for the past ten years. Her fa-

vorite exhibit is The Big Cat Exhibit. Over thousands of observations, the animal in this exhibit has never been a jaguar. Today Michelle must stay home and can't visit the zoo because she sprained her ankle. While relaxing on the couch, Michelle thinks, "The animal in the Big Cat Exhibit today is not a jaguar." And she is right: it is not a jaguar.

There was one difference in procedure between this experiment and Experiment 1: this time participants were also asked to rate their confidence in their answer to the knowledge question. Responses were collected on a 1 ("not at all confident") to 10 ("completely confident") scale. This was done to check for subtle differences of opinion not captured by the dichotomous measure. Otherwise, the questions and procedure for this experiment are the same as for Experiment 1.

3.2. Results

I made three predictions about the rate of knowledge denial: (1) it would vary across the conditions, (2) it would be lower in Positive Jaguar than in each of the other two conditions, (3) it would not differ between Negative Jaguar and Negative Leopard.

All three predictions were true. First, the rate of knowledge denial varied across the conditions, $\chi^2(2, N = 305) = 8.31$, $p = .016$, Cramer's V = .165. Second, pairwise comparisons revealed that the rate of knowledge denial was lower in Positive Jaguar (N = 100, 53%) than in Negative Jaguar (N = 103, 71.8%), Fisher's exact test, $p = .006$, Cramer's V = .195, and lower in Positive Jaguar than in Negative Leopard (N = 102, 66.7%), Fisher's exact test, $p = .033$, one-tailed, Cramer's V = .139. Third, rate of knowledge denial didn't differ between Negative Jaguar and Negative Leopard, Fisher's exact test, $p = .452$, n.s. Figure 2 visualizes these results.

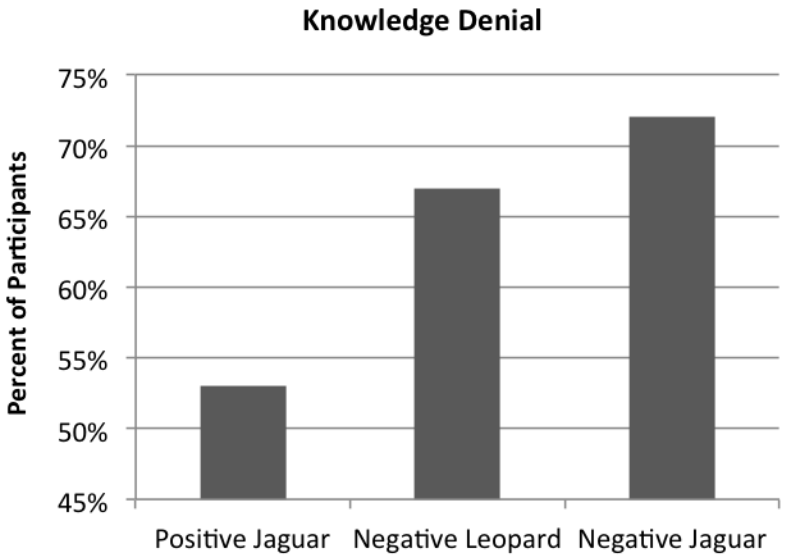

Fig. 2. Experiment 2. Denial of knowledge across conditions.

Binomial tests revealed that rate of knowledge denial was significantly above chance in both Negative Jaguar and Negative Leopard (*p*s < .001). By contrast, it was no different from chance in Positive Jaguar ($p = .617$).

Participants were also asked to rate their confidence in their answer to the knowledge question. A one-way analysis of variance detected no difference in confidence scores across the conditions: Positive Jaguar, $M = 8.67$, $SD = 1.68$; Negative Jaguar, $M = 8.77$, $SD = 1.34$; Negative Leopard, $M = 8.51$, $SD = 1.60$; $F(2) = .715$, $p = .49$, n.s.

3.3. Discussion

Two main things emerge from this study. First, the results replicate a main finding from Experiment 1: negative inferential belief is evaluated more harshly than positive inferential belief, even when each is based on a long and unbroken track record of evidence. Second, the results support the hypothesis that the skeptic exploits an evaluative bias against negative inferential beliefs. The skeptic's tricky

alternative scenarios aren't necessarily especially difficult to rule out. For instance, in this study participants were just as unwilling to ascribe knowledge to Michelle that the animal was not a jaguar (~O) as they were to ascribe knowledge that the animal was not a leopard (~T). By contrast, participants were much more willing to ascribe knowledge that the animal was a jaguar (O). In short, the key difference is not be between O and T but rather between negative and affirmative contents.

Once again it seems that the skeptic succeeds in getting us to seriously doubt ourselves only in cases of negative inferential belief.

4. Conclusion

Certain skeptical arguments have an uncanny ability to simultaneously attract and repel our intellects. They use premises that we find plausible, only to end with conclusions that we find absurd. What explains our ambivalent relationship to these arguments? Given that nearly no one actually accepts the skeptical conclusions, why do we pay any attention to them? Why is skepticism, despite its implausibility, a staple of both intellectual and popular culture?

Philosophers have proposed many explanations. Perhaps it's because we want to know things and understand what knowledge requires, but we implicitly recognize that we can't meet those requirements (Unger, 1975). Or perhaps there are deep tensions in our concept of knowledge, or in our pursuit of it, which lead us to vacillate on whether we can ever know anything (Stroud, 1984; Nagel, 1986: ch. 5). Or perhaps we misunderstand the relationship between the things we ordinarily know and the things the skeptic says we can't know: maybe we don't need to rule out skeptical alternatives in order to know ordinary things about the world around us (Dretske, 1970; Nozick, 1981). Or perhaps the verb 'knows' is semantically context-sensitive and the skeptic is able to manipulate the conversational

context so that he speaks truthfully when he says we "don't know" something (Cohen, 1988; DeRose, 1995; Lewis, 1996). Or perhaps the skeptic exploits facts about the pragmatics of knowledge discourse or the social function of knowledge attributions (Turri, 2010; Turri, 2012; Turri, in press).

But before drawing sweeping conclusions about the nature of knowledge or proposing revisionary semantic hypotheses, we might do well to identify psychological factors that contribute to skepticism's enduring appeal. I've developed a version of this last suggestion in this paper.

Our findings from two studies support several conclusions. First, the classic skeptical argument poses a greater threat to inferential belief than perceptual belief. Second, inferential belief with a negative conclusion is judged especially harshly and is unlikely to be classified as knowledge. Third, it turns out that classic skeptical arguments owe a good deal of their potency to this evaluative bias against negative inferential beliefs, namely, *the source-content bias*. Before the skeptic even shows up, people confidently claim that such beliefs don't count as knowledge. When the skeptic finally does show up to make his case, it probably does not, as some philosophers claim, "come as a surprise" — it is not "startling" or "shocking" to hear — that such beliefs don't seem to be knowledge. The skeptic has not "pointed out that we do not know things we would have confidently said we knew" (Nozick 1981: 202). People routinely deny that such beliefs count as knowledge. It is not, in the first place, some deep fact about the nature of knowledge or the skeptic's ingenuity that invites skepticism. It's our psychology.

The bias against negative content for inferential belief *specifically* is surprising and further work is required to better understand it. Voluminous behavioral research shows that it is more complicated and effortful to represent negation than affirmation (Wason, 1961; Wason, 1965; Gough, 1966; Slobin, 1966; MacDonald & Just, 1989; Evans, Newstead & Byrne, 1993; Kaup, 2001; Hasson & Glucksberg, 2006; Stupple & Waterhouse, 2009), and recent work in cognitive

neuroscience has begun uncovering the neural signatures of this difference (Christensen, 2009; Kumar, Padakannaya, Mishra & Khetrapal, 2013). Accordingly, one might expect knowledge ascriptions for positive claims to be higher in general, whether the source of belief is inference or perception, because "any variable that increases experienced ease of processing is also likely to increase judgments of truth" (Reber & Schwarz, 1999: 342). But this is not what we observed. Instead, we observed significantly lower rates of knowledge attribution for negative inferential beliefs but no corresponding asymmetry for perceptual beliefs.

Recent research suggests that a readily accessible schema for accommodating negative information promotes greater fluency (Mayo, Schul & Burnstein, 2004). For example, "it is not warm" can be more readily processed than "it is not wary" because the familiar "cold" accommodates "not warm," whereas no common adjective as readily accommodates "not wary." However, in the studies reported here, the protagonist's negative belief had the same content whether it was formed perceptually or inferentially, so a more accessible opposing conceptual schema can't explain the asymmetry. Thus it's possible that *the source of information* also facilitates comprehension of negative information. Perhaps people more fluently process negative information if the informational source is perceptual, thereby raising rates of knowledge ascription. In other words, in addition to features of informational *content* enhancing fluency, perhaps features of informational *source* can too.

This last suggestion gains some plausibility from a growing body of work on the perceptual simulation theory of linguistic comprehension. According to this view, when comprehending a text, people construct a mental simulation of the situation it describes (Johnson-Laird, 1983; Graesser, Millis & Zwaan, 1997) and this mental simulation shares the representational format of perception and action (Zwaan, 2004; Bianchi, Savardi, Burro & Torquati, 2011; see also Barsalou,

1999). Interestingly, one line of evidence for this view comes from a series of studies on *negation* (Kaup, 2001; Kaup & Zwaan, 2003; Kaup, Yaxley, Madden, Zwaan & Lüdtke, 2007; Zwaan & Percher, 2012 replicates relevant results using Amazon Mechanical Turk). If people comprehend negations by perceptually simulating them, then this might help explain why negation affects the evaluation of inferential belief differently than perceptual belief. For when the protagonist in the story forms her belief perceptually, her mode of cognition closely matches the mode in which participants simulate and thereby comprehend the case. By contrast, when the protagonist forms her belief non-perceptually, her mode of cognition mismatches the mode in which participants simulate the case. Participants asked to evaluate a negative inferential belief are thus required to dissociate two modes of cognition, their own and the protagonist's, resulting in a more complicated theory-of-mind task. This in turn could inhibit fluency and thus depress rates of knowledge attribution.

Regardless of whether this suggestion explains the results observed above, the experiments reported here provide evidence for two main conclusions. First, there is a source-content bias in cognitive evaluation and, second, classic forms of skepticism prey on this bias.

Acknowledgments — For helpful feedback and conversations, I'm happy to thank Peter Blouw, Wesley Buckwalter, Ori Friedman, Joshua Knobe, Duncan Pritchard, Blake Myers-Schulz, Christina Starmans, David Rose, and Angelo Turri. Thanks also to the referees and editors of *Cognitive Science* and audiences at Rutgers University and the 2013 Buffalo Experimental Philosophy Conference. This research was kindly 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

Bartsch, K. (2002). The role of experience in children's developing folk epistemology: review and analysis from the theory–theory perspective. *New Ideas in Psychology*, *20*(2-3), 145–161.

Bayle, P. (1697 [1991]). *Historical and critical dictionary: selections*. (R.H. Popkin, Trans.). Indianapolis: Hackett.

Beebe, J. R., & Buckwalter, W. (2010). The Epistemic Side-Effect Effect. *Mind & Language*, *25*(4), 474–498.

Beebe, J. R., & Jensen, M. (2012). Surprising connections between knowledge and action : The robustness of the epistemic side-effect effect. *Philosophical Psychology*, *25*(5), 689–715.

Berkeley, G. (1710 [1982]). *Treatise concerning the principles of human knowledge*. (K. Winkler, Ed.). Indianapolis: Hackett.

Bianchi, I., Savardi, U., Burro, R., & Torquati, S. (2011). Negation and psychological dimensions. Journal of Cognitive Psychology, 23(3), 275–301.

Buckwalter, W., & Schaffer, J. Forthcoming. Knowledge, stakes, and mistakes. *Noûs*.

Buckwalter, W., Rose, D., & Turri, J. In press. Belief through thick and thin. *Noûs*.

Buckwalter, W. & Turri, J. (Under review). In the thick of moral motivation.

Call, J., & Tomasello, M. (2008). Does the chimpanzee have a theory of mind? 30 years later. Trends in Cognitive Sciences.

Christensen, K. R. (2009). Negative and affirmative sentences increase activation in different areas in the brain. Journal of Neurolinguistics, 22(1), 1–17.

Cohen, S. (1988). How to be a Fallibilist. *Philosophical Perspectives*, *2*, 91–123.

Cohen, S. (2005). Contextualism defended. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (1st ed., pp. 56–62). Blackwell.

Cushman, F., & Greene, J. D. (2012). Finding faults: How moral dilemmas illuminate cognitive structure. *Social Neuroscience*, 7(3), 269–279.

Dally, J. M., Emery, N. J., & Clayton, N. S. (2005). Cache protection strategies by western scrub-jays, Aphelocoma californica: implications for social cognition. *Animal Behaviour*, 70(6), 1251–1263.

DeRose, K. (1995). Solving the skeptical problem. *Philosophical Review*, *104*(1), 1–52.

Descartes, R. (1641). *Meditations on first philosophy*.

Dretske, F. (1970). Epistemic operators. *The Journal of Philosophy*, *67*(24), 1007–1023.

Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, Mass.: MIT Press.

Feltz, A., & Zarpentine, C. (2010). Do you know more when it matters less? *Philosophical Psychology*, 23(5), 683–706.

Flavell, J. H. (2000). Development of children's knowledge about the mental world. *International Journal of Behavioral Development*, *24*(1), 15–23.

Friedman, O. & Turri, J. (Under review). Is probabilistic evidence a source of knowledge?

Goldman, A. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, *73*(20), 771–791.

Goldman, A. (1993). *Philosophical applications of cognitive science*. Boulder: Westview Press.

Gough, P. B. (1966). The verification of sentences: the effects of delay of evidence and sentence length. *Journal of verbal learning and verbal behavior*, *5*, 492–496.

Graesser, A.C., Millis, K.K. & Zwaan, R.A. (1997). Discourse comprehension. *Annual review of psychology*, *48*, 163–189.

Greco, J. (2000). *Putting skeptics in their place: the nature of skeptical argu-*

ments and their role in philosophical inquiry. Cambridge: Cambridge University Press.

Greene, J. D., Sommerville, R. B., Nystrom, L. E., & Darley, J. M. (2001). An fMRI Investigation of Emotional Engagement in Moral Judgment. *Science*, 293, 2105–2108.

Hardy-vallée, B., & Dubreuil, B. (2010). Folk Epistemology as Normative Social Cognition, (December 2009), 483–498. doi:10.1007/s13164-009-0020-5.

Hare, B., Call, J., & Tomasello, M. (2001). Do chimpanzees know what conspecifics know? *Animal behaviour*, *61*(1), 139–151.

Hasson, U., & Glucksberg, S. (2006). Does understanding negation entail affirmation? *Journal of Pragmatics*, *38*(7), 1015–1032. doi:10.1016/j.pragma.2005.12.005

Horowitz, A. (2008). Attention to attention in domestic dog (Canis familiaris) dyadic play. *Animal Cognition*, 12(1), 107–118.

Hume, D. (1739). *A treatise of human nature*. L. A. Selby-Bigge & P. H. Nidditch (Eds.) (2nd ed., 1979). Oxford: Oxford University Press.

Hume, D. (1748). *An enquiry concerning human understanding*. Eric Steinberg (Ed.). Indianapolis: Hackett, 1993.

Johnson-Laird, P.N. (1983). *Mental models*. Cambridge, MA: Harvard University Press.

Kaup, B. (2001). Negation and its impact on the accessibility of text information. *Memory & Cognition, 29*(7), 960–967.

Kaup, B., & Zwaan, R. A. (2003). Effects of negation and situational presence on the accessibility of text information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(3), 439–446.

Kaup, B., Yaxley, R. H., Madden, C. J., Zwaan, R. A., & Lüdtke, J. (2007). Experiential simulations of negated text information. *The Quarterly Journal of Experimental Psychology, 60*(7), 976–990.

Keefner, A. 2013. *Do birds have a theory of mind?* (Master's thesis.) University of Waterloo.

Kumar, U., Padakannaya, P., Mishra, R. K., & Khetrapal, C. L. (2013). Distinctive neural signatures for negative sentences in Hindi: an fMRI study. *Brain Imaging and Behavior*, 7(2), 91–101.

Mayo, R., Schul, Y., & Burnstein, E. (2004). "I am not guilty" vs 'I am innocent': Successful negation may depend on the schema used for its encoding. *Journal of Experimental Social Psychology*, 40(4), 433–449.

Mercier, H. (2010). The Social Origins of Folk Epistemology. *Review of Philosophy and Psychology*, *1*(4), 499–514.

Lewis, D. (1979). Scorekeeping in a language game. *Journal of Philosophical Logic*, *8*(1), 339–359.

Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, *74*(4), 549–567.

MacDonald, M. C., & Just, M. A. (2002). Changes in activation levels with negation. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*(4), 633–642.

Mayo, R., Schul, Y., & Burnstein, E. (2004). "I am not guilty" vs 'I am innocent': Successful negation may depend on the schema used for its encoding. *Journal of Experimental Social Psychology*, 40(4), 433–449.

Kumar, U., Padakannaya, P., Mishra, R. K., & Khetrapal, C. L. (2012). Distinctive neural signatures for negative sentences in Hindi: an fMRI study. Brain Imaging and Behavior.

Moore, G.E. (1959). *Philosophical papers*. New York: Collier.

Murray, D., Sytsma, J., & Livengood, J. (in press). God knows (but does God believe?). Philosophical Studies. doi:10.1007/s11098-012-0022-5

Myers-Schulz, B., & Schwitzgebel, E. (2013). Knowing that P without believing that P. *Noûs*, 47(2), 371–384.

Nagel, J., San Juan, V., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. Cognition, 1–10. doi:10.1016/j.cognition.2013.02.008

Nagel, T. (1986). *The view from nowhere*. Oxford: Oxford University Press.

Nozick, R. (1981). *Philosophical explanations*. Harvard University Press.

Oppenheimer, D. M. (2008). The secret life of fluency. *Trends in Cognitive Sciences*, 12(6), 237–241.

Perner, J., Frith, U., Leslie, A. M., & Leekam, S. R. (1989). Exploration of the autistic child's theory of mind: Knowledge, belief, and communication. Child development, 60(3), 689–700.

Pillow, B. H. (1989). Early understanding of perception as a source of knowledge. *Journal of experimental child psychology*, *47*(1), 116–29.

Povinelli, D., Nelson, K., & Boysen, S. (1990). Inferences about guessing and knowing by chimpanzees. *Journal of Comparative Psychology*, *104*(3), 203–210.

Povinelli, D., & DeBlois, S. (1992). Young children's understanding of knowledge formation in themselves and others. *Journal of Comparative Psychology*, *106*(3), 228–238.

Pritchard, D. (2002). Recent work on radical skepticism. *American Philosophical Quarterly*, *39*(3), 215–257.

Reber, R., & Schwarz, N. (1999). Effects of perceptual fluency on judgments of truth. *Consciousness and Cognition, 8*(3), 338–342.

Reid, T. (1764). *An inquiry into the human mind on principles of common sense*. D. R. Brookes (Ed.). University Park, Penn.: Pennsylvania State University Press.

Reiss, D. (2001). Mirror self-recognition in the bottlenose dolphin: A case of cognitive convergence. P*roceedings of the National Academy of Sciences*, 98(10), 5937–5942.

Russell, B. (1921). *The analysis of mind*. Project Gutenberg. Retrieved from

http://www.gutenberg.org/files/2529/2529-h/2529-h.htm.

Schaffer, J., & Knobe, J. (2012). Contrastive Knowledge Surveyed. *Noûs*, *46*(4), 675–708.

Sellars, W. (1963). *Science, perception and reality*. Atascadero, Calif.: Ridgeview Publishing Company.

Sextus Empiricus. *Outlines of Pyrrhonism*.

Slobin, D.I. (1966). Grammatical transformations and sentence comprehension in childhood and adulthood. *Journal of Verbal Learning and Verbal Behavior*, *5*, 219–227.

Sodian, B., & Wimmer, H. (1987). Children's understanding of inference as a source of knowledge. *Child development*, 58(2), 424–433.

Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Topics*, *13*, 141–153.

Spicer, F. (2010). Cultural Variations in Folk Epistemic Intuitions. *Review of Philosophy and Psychology*, *1*(4), 515–529.

Sprung, M. (2010). Clinically Relevant Measures of Children's Theory of Mind and Knowledge about Thinking: Non-Standard and Advanced Measures. *Child and Adolescent Mental Health*, *15*(4), 204–216.

Sripada, C. S., & Stanley, J. (2012). Empirical tests of interest-relative invariantism. *Episteme*, 9(1), 3–26.

Starmans, C., & Friedman, O. (2012). The folk conception of knowledge. *Cognition*, *124*(3), 272–83.

Stine, G. (1976). Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, *29*, 249–261.

Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and social psychology review 8*(3), 220–47.

Stroud, B. (1984). *The significance of philosophical skepticism*. Oxford: Clarendon Press.

Stupple, E. J. N., & Waterhouse, E. F. (2009). Negations in syllogistic reasoning: Evidence for a heuristic–analytic conflict. *The Quarterly Journal of Experimental Psychology*, *62*(8), 1533–1541.

Swain, S., Alexander, J., & Weinberg, J. M. (2008). The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp. *Philosophy and Phenomenological Research*, 84(1), 138–155.

Turri, J. (2010). Epistemic Invariantism and Speech Act Contextualism. Philosophical Review, 119(1), 77–95.

Turri, J. (2012). Pyrrhonian skepticism meets speech-act theory. *International Journal for the Study of Skepticism*, *2*(1), 83–98.

Turri, J. (2013). A conspicuous art: putting Gettier to the test. *Philosophers' imprint*.

Turri, J. (2014). *Epistemology: a guide*. Malden, Mass.: Wiley-Blackwell.

Turri, J. (In press). Linguistic intuitions in context: a defense of nonskeptical pure invariantism. In A. Booth & D. Rowbottom (Eds.), *Intuitions*. Oxford University Press.

Turri, J. (Under review). An open and shut case: epistemic closure in the manifest image.

Turri, J. & Friedman, O. (in press). Winners and losers in the folk epistemology of lotteries. In Beebe, J. (Ed.), *Advances in experimental epistemology*. Continuum.

Turri, J. & Friedman, O. (under review). A peculiar and perpetual error: An asymmetry in judgments about the detection of affirmative and negative outcomes.

Unger, P. (1975). *Ignorance: a case for scepticism*. Oxford: Oxford University Press.

Vogel, J. (1990). Are there counterexamples to the closure principle? In M. D. Roth & G. Ross (Eds.), *Doubting* (pp. 13–27). Dordrecht: Kluwer Aca-

demic Publishers. Reprinted in *Epistemology: an anthology,* 2ed. Ed. E. Sosa, J. Kim, J. Fantl & M. McGrath. Blackwell.

Watson, P. C. (1961). Response to affirmative and negative binary statements. *British Journal of Psychology, 52*(2), 133–142.

Wason, P. C. (1965). The contexts of plausible denial. *Journal of verbal learning and verbal behavior*, *4*, 7–11.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

Williamson, T. (2006). Can cognition be factored into internal and external components? In R. Stainton (Ed.), *Contemporary debates in cognitive science*, pp. 291–306. Blackwell.

Wright, J. C. (2010). On intuitional stability: the clear, the strong, and the paradigmatic. *Cognition*, *115*(3), 491–503.

Zamir, E., Ritov, I., & Teichman, D. (In press). Seeing is Believing: The Anti-Inference Bias. *Indiana Law Journal*. doi:10.2139/ssrn.1989561, accessed August 2013.

Zwaan, R.A. (2004). The immersed experiencer: toward an embodied theory of language comprehension. In B.H. Ross (Ed.), *The psychology of learning and motivation* (vol. 44, pp. 35–62). New York: Academic Press.

Zwaan, R. A., & Pecher, D. (2012). Revisiting mental simulation in language comprehension: six replication attempts. *PLOS One*, 7(12), 1–10.