Actionability Judgments Cause Knowledge Judgments*

John Turri

Wesley Buckwalter

David Rose

> The germinal question concerning things brought for the first time before consciousness is . . . the practical 'What is to be done?'.
>
> — William James (1879, p. 18)

Abstract: Researchers recently demonstrated a strong direct relationship between judgments about what a person knows ("knowledge judgments") and judgments about how a person should act ("actionability judgments"). But it remains unknown whether actionability judgments cause knowledge judgments, or knowledge judgments cause actionability judgments. This paper uses causal modeling to help answer this question. Across two experiments, we found evidence that actionability judgments cause knowledge judgments.

Keywords: knowledge; action; pragmatism; stakes; social cognition

Introduction

A foundational debate in philosophy concerns the relationship between knowledge and action. Rene Descartes claimed that knowledge is fundamentally separable from action and other practical concerns. Descartes cited this assumption in defense of his infamous method of hyperbolic doubt, introduced in the First Meditation (Descartes, 1641). John Locke responded to

* This is the penultimate version of a paper forthcoming in *Thought*. Please cite the final, published version if possible.

professed Cartesian doubt by rejecting the assumption that knowledge is fundamentally separable from action. Locke claimed that knowledge was connected to whether you have "assurance enough" to "govern [your] actions" (Locke 1690, book 4.11.8). William James went one step further than Locke. James claimed that knowledge is partly *constituted* by "practical interests" and that it was essentially "a guide to appropriate action" (James, 1879, p. 18). The disagreement between Descartes and Locke was forgotten for most of the twentieth century in Anglo-American philosophy, but it has recently been renewed, with many epistemologists adopting Cartesian positions (e.g. Fumerton, 2010; Reed, 2013) and others defending Lockean positions (e.g. Fantl & McGrath 2002, 2009; Hawthorne, 2004; Stanley, 2005; Hawthorne & Stanley, 2008).

In his response to Descartes, Locke suggested that observations about how we *talk* and *act* are relevant to a correct theoretical understanding of the relationship between knowledge and practice (Locke 1690/1975, bk. 4.11.3-8). The emphasis on behavioral evidence plays a crucial role in contemporary defenses of Lockeanism about knowledge (e.g. Fantl & McGrath, 2007, p. 562; Hawthorne & Stanley, 2008, p. 571). In particular, latter-day Lockeans claim that ordinary knowledge attributions are sensitive to practical factors, such as an increase in stakes or a change in how you should act. Cartesians have largely conceded Lockean claims about our behavior, but they propose that the sensitivity is entirely indirect (e.g. Weatherson, 2005; Bach, 2005; Ganson, 2007; Nagel, 2008). More specifically, they argue that practical factors affect knowledge attributions only by influencing our estimation of the factors that, according to Cartesians, constitute knowledge: belief, truth, and quality of evidence.

Until recently, the philosophical debate has relied mainly on introspection and social

observation to evaluate the behavioral claims, which is typical in Anglophone analytic philosophy (for reviews, see Turri, 2016a; Turri, 2016b; Turri 2016c). But introspection and social observation are subject to well known limitations (e.g. Milgram, 1974, pp. 103-4; Nisbett & Wilson, 1977; Wilson, 2002; Lieberman, 2013), and early experimental attempts to validate the behavioral claims met with mixed results (e.g. Feltz & Zarpentine, 2010; May, Sinnott-Armstrong, Hull & Zimmerman, 2010; Buckwalter, 2010; Pinillos, 2012; Sripada & Stanley, 2012). As two commentators recently put it, the evidence from these early attempts "made it easy" to conclude that much of the theoretical debate "was founded on a myth" (Schaffer & Knobe, 2012, p. 675). Equally worrisome, moreover, is the fact that Cartesians and Lockeans propose competing psychological models of knowledge judgments that introspection and social observation cannot adjudicate. Even supposing that judgments about how someone should act ("actionability judgments") are connected to knowledge judgments, we cannot introspect or directly observe whether this relationship is direct, as latter-day Lockeans claim, or indirect, as latter-day Cartesians claim.

A recent study addressed this issue (Turri & Buckwalter, in press). Using regression and mediation analysis, researchers found a strong direct relationship between actionability judgments and knowledge judgments. Indeed, this relationship was stronger than the relationship between knowledge judgments and evidence judgments, and it was approximately as strong as the relationship between knowledge judgments and truth judgments. These findings rule out Cartesian mediation models and fit very well with Lockean models. But they leave an important question unanswered. Given that there is a strong direct relationship between actionability judgments and knowledge judgments, what is the nature of that link? Which judgment, if either,

has priority? Do knowledge judgments cause actionability judgments, do actionability judgments cause knowledge judgments, or do both have a common cause?

This paper reports two experiments designed to answer this question. The experiments are modelled after earlier work that ruled out Cartesian models. But they go beyond earlier work in two crucial ways.

First, they are more closely informed by relevant work in the judgment and decision making literature. In particular, they are informed by important work that identifies three separate factors that contribute to the perception of heightened stakes (McAllister, Mitchell & Beach, 1979; Beach & Mitchell, 1976). The three factors are irreversibility, significance and accountability. An agent's conduct is irreversible to the extent that it cannot be undone to remove its negative consequences. The more irreversible conduct is, the more likely people are to perceive the situation as high stakes. An agent's conduct is significant to the extent that its likely consequences are extremely good, extremely bad, or have an extremely broad impact. The more significant conduct is, the more likely people are to perceive the situation as high stakes. An agent is accountable for conduct to the extent that she is responsible for its consequences. She will be praised or blamed, rewarded or punished, depending on whether the outcome is positive or negative. The more accountable an agent is for her conduct, the more likely people are to perceive the situation as high stakes. In the experiments that follow, a "high stakes" case is one that is high on all three dimensions of stakes, and a "low stakes" case is one that is low on all three dimensions.

Second, the present research uses a causal search with the Greedy Equivalence Search (GES) algorithm to identify the model that best fits the data. The causal search was conducted

with Tetrad 5.0.0 (http://www.phil.cmu.edu/projects/tetrad/). GES considers all possible models available given the different variables. Each variable is treated as a node. GES assigns an information score to the model in which all the nodes are disconnected — the "null model." GES then evaluates adding causal arrows — "edges" — between the nodes (Meek, 1997 provides the edge orientation rules). GES adds edges that best improve the model's information score, if such edges exist, until adding more edges does not improve the information score. At this point, GES evaluates whether deleting any edges will further improve the information score, and it deletes any such edges until deleting more edges does not improve the information score. In this way, GES helps us to identify the underlying causal structure among a set of variables, going beyond what regression and mediation analyses can offer. Regression assumes a causal direction, whereas GES does not, and GES is preferable to mediation models because it provides an overall measure of model fit and tends to deliver more accurate models (Iacobucci, Saldanha & Deng, 2007). GES returns the causal model that best fits the data and, given enough data, will return the true causal model (Chickering, 2002). (For other applications of GES to theoretical debates, see Rose & Nichols, 2013; Rose, Livengood, Sytsma & Machery, 2012; Turri 2016d).

Experiment 1

Method

Participants. Two hundred participants (aged 18–68 years, mean age = 31 years; 94% reporting English as a native language; 80 female) were tested. Participants were recruited and tested online using Amazon Mechanical Turk and Qualtrics and compensated $0.40 for approximately 2–3 minutes of their time. Repeat participation was prevented.

Materials and Procedure. Participants were randomly assigned to one of two conditions,

Low and High, in a between-subjects design. Each participant read a single story about an intelligence agent developing a file on a foreign operative, Ivan. Participants in the Low condition read a story that featured low irreversibility, significance and accountability. Participants in the High condition read a story that featured high irreversibility, significance, and accountability. The stories for the two conditions thus differed on all three crucial dimensions of stakes identified in the judgment and decision making literature. Here is the text for the stories, with manipulations in brackets and separated by a slash (Low/High):

> Jennifer is an intelligence analyst developing a file on Ivan, an elusive foreign operative. Jennifer has a source who tells her that Ivan stopped [his low-carb diet/selling arms to terrorists] and is no longer [jogging regularly/a threat]. ¶ Jennifer must submit a [provisional/final] report on Ivan to her supervisor within the hour. She will [definitely/definitely not] have a chance to revise her [provisional/final] report, and she [will not/will] be held accountable for decisions based on her [provisional/final] report.

After reading the story, participants were asked to rate their agreement or disagreement with eight statements.

1. If Jennifer changes her mind, she can revise her report later.
2. If Ivan still [jogs regularly/is a threat], there will be serious consequences.
3. Jennifer is responsible for decisions made based on her [provisional/final] report.
4. Jennifer thinks that Ivan no longer [jogs regularly/is a threat].
5. It's true that Ivan no longer [jogs regularly/is a threat].
6. Jennifer has good evidence that Ivan no longer [jogs regularly/is a threat].

7. Jennifer should write in the report that Ivan no longer [jogs regularly/is a threat].
8. Jennifer knows that Ivan no longer [jogs regularly/is a threat].

Statements 1–3 are manipulation checks on the three critical dimensions of stakes: irreversibility, significance and accountability. Statements 4–8 are the key dependent variables; responses to these statements will count as the participant's belief score, truth score, evidence score, actionability score, and knowledge score, respectively.

Responses were collected on a standard seven-point Likert scale anchored with "Strongly Disagree," "Disagree," "Somewhat Disagree," "Neutral," "Somewhat Agree," "Agree," and "Strongly Agree," left-to-right on the participant's screen. Responses were coded 1 (Strongly Disagree) to 7 (Strongly Agree). Participants never saw the numerical values, only the qualitative anchors. The eight statements were presented in random order and appeared on the participant's screen all at once, while the story remained at the top of the screen. Response options were always presented in the same order. After rating the statements, participants proceeded to a new screen where they completed a brief demographic questionnaire.

Results

Preliminary analysis revealed no main or interaction effects of participant age and gender on any of the eight dependent variables. Preliminary analysis also revealed that the stakes manipulation was extremely effective in all three dimensions (see Table 1).

Table 1. Experiment 1: Mean scores (SD) for the dependent measures in the Low and High conditions along with the results from independent samples t-tests.

Measure	Low	High	t	df	p	MD	d	95% CI for MD	
								LLCI	ULCI
Reversibility	6.03 (1.33)	1.71 (1.24)	23.67	198	<.001	4.32	3.36	3.96	4.67

Significance	2.91 (1.6)	6.22 (0.87)	-18.09	150	<.001	-3.31	2.56	-3.67	-2.95
Accountability	3.60 (2.09)	6.30 (0.90)	-11.81	132	<.001	-2.70	2.07	-3.15	-2.25
Belief	5.82 (1.15)	4.61 (1.30)	6.97	196	<.001	1.20	1.00	0.86	1.55
True	4.88 (1.04)	3.91 (1.35)	5.67	198	<.001	0.97	0.81	0.63	1.30
Evidence	4.91 (1.30)	4.11 (1.54)	3.97	194	<.001	0.80	0.57	0.40	1.20
Actionability	5.20 (1.46)	3.60 (1.69)	7.16	195	<.001	1.60	1.02	1.16	2.04
Knowledge	4.53 (1.73)	3.49 (1.65)	4.36	198	<.001	1.04	0.62	0.57	1.51

The independent variable Stakes (Low/High) and response to the dependent measures were entered into a causal search using the Greedy Equivalence Search algorithm in Tetrad V (see the Introduction). The model was constrained so that the independent variable Stakes could not be caused by any other variable in the model (reflecting the experimental design), and the penalty discount was set to 0.5 (decided *a priori*). Inspection of normal probability plots revealed no violation of normality assumptions. In line with previous research using similar methodology (Turri & Buckwalter, in press), preliminary regression analysis revealed that belief scores did not predict either knowledge scores or actionability scores. GES also revealed that including belief scores resulted in a poor fitting model, so we omitted them.[1] Figure 1 depicts the best fitting model, which fits the data well, $\chi^2(3) = 3.44$, $p = .33$, BIC = -12.45. In the model, actionability judgments cause knowledge judgments. Indeed, actionability judgments alone directly cause knowledge judgments, and this is the strongest causal relationship between any two dependent variables in the model.

1 It is perhaps worth noting that the best fitting model that included belief featured an edge from actionability judgments to knowledge judgments. That is, including belief resulted in the same causal relationship between actionability and knowledge judgments reported below.

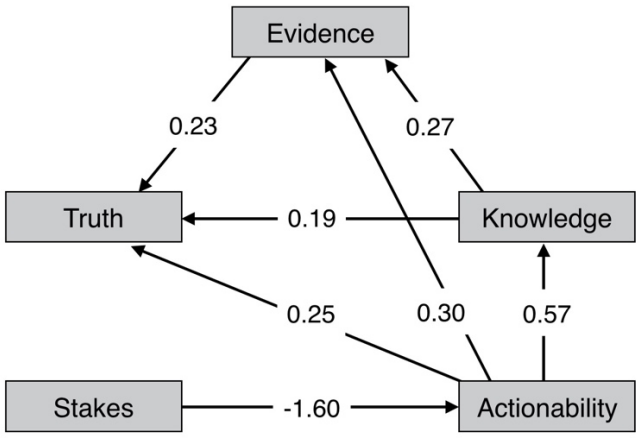

Figure 1. Experiment 1. Graph of the best fitting model as determined by a causal search using the Greedy Equivalence Search algorithm. Arrows represent directional causal pathways. Path coefficients represent the causal relation's strength and quality. Positive numbers represent a positive causal relation (higher scores for the cause generate higher scores for the effect; lower scores for the cause generate lower scores for the effect); negative numbers represent a negative causal relation (higher scores for the cause generate lower scores for the effect; lower scores for the cause generate higher scores for the effect).

As a point of comparison, we conducted the causal search in the exact same way except that it forced an edge from knowledge judgments to actionability judgments. This model was rejected as a poor fit for the data, $\chi^2(5) = 44.75$, $p = 0$, BIC = 18.27.

Discussion

Our main question was to identity the direction of the relationship between knowledge judgments and actionability judgments. In the present study, actionability judgments caused knowledge judgments and this was the strongest causal relationship between any two dependent variables. The next experiment investigates the robustness of this finding by testing a different cover story and probing for attributions of knowledge and actionability only.

Experiment 2

Method

Participants. Two hundred and five participants (aged 18–72 years, mean age = 32 years; 96% reporting English as a native language; 101 female) were tested. Participants were recruited and tested the same way as in Experiment 1. Repeat participation was prevented.

Materials and Procedure. Participants were randomly assigned to one of two conditions, Low and High, in a between-subjects design. Each participant read a single story about an employee charged with updating the menu for a chain of coffee shops. As in Experiment 1, the stories for the two conditions thus differed on all three crucial dimensions of stakes identified in the judgment and decision making literature. Here is the text for the stories, with manipulations in brackets and separated by a slash (Low/High):

> Christina is in charge of updating the menu for a chain of coffee shops each day. To some customers [interested in the history and culture of coffee/with severe nut allergies], it matters whether the coffee [is from northern Colombia/contains pine nuts]. While working on today's menu, Christina observes that the latest shipment of coffee [is from northern Colombia/contains trace amounts of pine nuts]. ¶ Christina must distribute a [provisional/final] copy of the menu to all the shops in just a couple minutes. She [will/will not] have a chance to revise the menu before the shops open, and she [will not/will] be held accountable for the accuracy of the menu.

After reading the story, participants were asked to rate their agreement or disagreement with five statements.

1. Christina will be held accountable for the accuracy of the menu.
2. Christina will have a chance to revise the menu.

3. If the coffee [is from northern Colombia/contains pine nuts], there could be serious consequences.
4. Christina should write on the menu that the coffee [is from northern Colombia/contains pine nuts].
5. Christina knows that the coffee [is from northern Colombia/contains pine nuts].

Statements 1–3 are manipulation checks on the three critical dimensions of stakes. Statements 4 and 5 are the key dependent variables of knowledge and actionability. Responses were collected and scored the same way as in Experiment 1, including randomized order for all statements.

Results

Preliminary analysis revealed no main or interaction effects of participant age and gender on any of the dependent variables. Preliminary analysis also revealed that the stakes manipulation was extremely effective in all three dimensions (see Table 2).

Table 2. Experiment 2. Mean scores (SD) for the dependent measures in the Low and High conditions along with the results from independent samples t-tests.

Measure	Low	High	t	df	p	MD	d	95% C.I. for MD	
								LLCI	ULCI
Reversibility	6.15 (1.28)	2.43 (1.88)	16.56	178	<.001	3.71	2.48	3.27	4.16
Significance	2.86 (1.63)	6.29 (1.13)	-17.50	182	<.001	-3.43	2.59	-3.82	-3.04
Accountability	2.42 (1.99)	6.44 (0.91)	-18.67	143	<.001	-4.02	3.12	-4.45	-3.60
Actionability	6.00 (0.97)	6.56 (0.77)	-4.58	203	<.001	-0.56	0.64	-0.80	-0.32
Knowledge	6.13 (1.11)	6.19 (1.13)	-0.38	203	.702	-0.06	0.05	-0.37	0.25

We conducted a causal search on the data in the same way as Experiment 1. Figure 2 depicts the best fitting model, which fits the data well, $\chi^2(1) = 2.68$, $p = .11$, BIC = -2.64. In the

model, actionability judgments cause knowledge judgments.

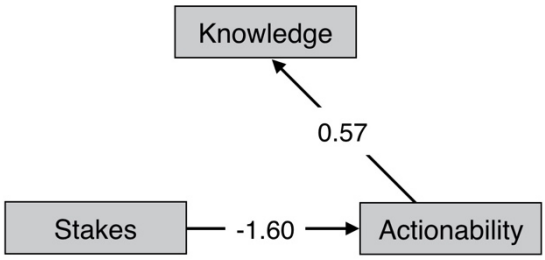

Figure 2. Experiment 2. Graph of the best fitting model as determined by a causal search using the Greedy Equivalence Search algorithm. Arrows represent directional causal pathways. Path coefficients represent the causal relation's strength and quality. Positive numbers represent a positive causal relation (higher scores for the cause generate higher scores for the effect; lower scores for the cause general lower scores for the effect); negative numbers represent a negative causal relation (higher scores for the cause generate lower scores for the effect; lower scores for the cause generate higher scores for the effect).

As a point of comparison, an alternative model in which stakes cause knowledge judgments and, in turn, knowledge judgments cause actionability judgments was rejected as a poor fit for the data, $\chi^2(1) = 22.56$, $p < .00001$, BIC = 17.24.

Discussion

The main finding from Experiment 1 was that, in the best fitting causal model, actionability judgments cause knowledge judgments. The present experiment replicated that finding and generalized it to other narrative contexts. It's also worth noting two further points. First, in Experiment 1 higher stakes caused lower actionability scores, whereas in the present experiment higher stakes caused higher actionability scores. This could be because the scenarios for Experiments 1 and 2 differ in whether the riskier option is to *act* or *not act* (for related results on

knowledge judgments and actionability judgments, see Turri, Friedman & Keefner, in press). Second, in Experiment 1 the stakes manipulation caused (indirectly) a significant difference in knowledge judgments, whereas in the present experiment stakes did not affect knowledge judgments. Despite these dissimilarities, actionability judgments caused knowledge judgments in both experiments.

Conclusion

Recent work has demonstrated a strong direct relationship between judgments about what someone knows ("knowledge judgments") and how they should act (i.e. "actionability judgments"). This result answered a longstanding question in philosophy about whether the relationship between these judgments is direct or instead mediated by other judgments about truth, belief and evidence. But it left an important question unanswered: do actionability judgments cause knowledge judgments, or do knowledge judgments cause actionability judgments? The present research was designed to help answer this question.

Across two experiments, we found evidence that judgments about how someone should act cause judgments about whether they know. We observed this same basic result in different narrative contexts, when higher stakes caused lower actionability judgments, and when higher stakes caused higher actionability judgments. These results do not show that actionability judgments always cause knowledge judgments, or that knowledge judgments never cause actionability judgments. Further work is required to test how broadly we observe the same basic causal pattern reported above. Nevertheless, the results do provide evidence that in a range of perfectly ordinary cases, actionability judgments cause knowledge judgments.

One assumption of our causal search analyses is that our model does not omit an important

latent variable, which might be a common cause of both actionability judgments and knowledge judgments. In Experiment 1, we measured a suite of judgments that theorists commonly assume inform judgments of knowledge or actionability. Our analysis suggests that even when accounting for these other judgments, actionability judgments cause knowledge judgments. Nevertheless, we cannot rule out the theoretical possibility that an unmeasured common cause drove some of the findings. Future work could aim to identify and measure such variables and test whether findings similar to ours persist in that context.

It is worth noting that the best fitting model for Experiment 1 contains some other potentially very interesting features. For instance, according to this model, the first judgment people make concerns actionability, which in turn influences a host of other judgments, including judgments of knowledge, evidence, and truth. Although this pattern is not predicted by any contemporary philosophical theory we are aware of, it was predicted by William James. Wrote James, "The germinal question concerning things brought for the first time before consciousness is . . . the practical 'What is to be done?'" (James, 1979, p. 18). To take another example, belief attributions did not affect knowledge judgments, whereas a mainstream view in epistemology is that belief is an "ingredient" of knowledge. Nevertheless, even if this result does not cohere with mainstream philosophical assumptions, it fits well with recent research on the psychology of knowledge attributions. This research found that people often attribute knowledge while denying belief, or do not base knowledge attributions on belief attributions (e.g. Myers-Schulz & Schwitzgebel, 2013; Murray, Sytsma & Livengood, 2012; Turri & Buckwalter, in press).

Acknowledgements — For helpful feedback, we thank Angelo Turri and this journal's

anonymous referees. This research was supported by the Social Sciences and Humanities Research Council of Canada, the Ontario Ministry of Economic Development and Innovation, a Banting Postdoctoral Fellowship, and the Canada Research Chairs program.

References

Bach, K. (2005). The emperor's new 'knows'. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: knowledge, meaning, and truth* (pp. 51–89). Oxford: Oxford University Press.

Beach, L. R., & Mitchell, T. R. (1978). A Contingency Model for the Selection of Decision Strategies. *Academy of Management Review*, 3(3), 439–449.

Buckwalter, W. (2010). Knowledge Isn't Closed on Saturday: A Study in Ordinary Language. *Review of Philosophy and Psychology*, 1(3), 395–406.

Chickering, D. M. (2002). Optimal structure identification with greedy search. *The Journal of Machine Learning Research*, 3(1), 507–554.

Descartes, R. (1641/2006). *Meditations on first philosophy*. In R. Ariew & D. Cress (Ed. & Trans.), *Meditations, objections, and replies*. Indianapolis: Hackett.

Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *The Philosophical Review*, 111(1), 67–94.

Fantl, J., & McGrath, M. (2007). On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research*, 75(3), 558–589.

Fantl, J., & McGrath, M. (2009). Advice for fallibilists: put knowledge to work. *Philosophical Studies*, 142(1), 55–66.

Feltz, A., & Zarpentine, C. (2010). Do you know more when it matters less? *Philosophical*

Psychology, 23(5), 683–706.

Fumerton, R. (2010). Fencing out pragmatic encroachment. *Philosophical Perspectives*, 24, 243–253.

Ganson, D. (2007). Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139(3), 441–458.

Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.

Hawthorne, J., & Stanley, J. (2008). Knowledge and action. *Journal of Philosophy*, 105(10), 571.

Iacobucci, D., Saldanha, N., & Deng, X. (2007). A meditation on mediation: evidence that structural equations models perform better than regressions. *Journal of Consumer Psychology*, 17(2), 139–153.

James, W. (1879). The sentiment of rationality. In A. Castell (Ed.), *Essays in pragmatism*. New York: Hafner Press.

Lieberman, M. D. (2013). *Social: why our brains are wired to connect*. New York: Crown Publishers.

Locke, J. (1690/1975). *An essay concerning human understanding*. (P. H. Nidditch, Ed.). Oxford: Clarendon Press.

May, J., Sinnott-Armstrong, W., Hull, J. G., & Zimmerman, A. (2010). Practical interests, relevant alternatives, and knowledge attributions: an empirical study. *Review of Philosophy and Psychology*, 1(2), 265–273.

McAllister, D. W., Mitchell, T. R., & Beach, L. R. (1979). The contingency model for the selection of decision strategies: An empirical test of the effects of significance, accountability, and reversibility. *Organizational Behavior and Human Performance*, 24,

228–244.

Milgram, S. (1974). *Obedience to authority: an experimental view*. New York: Harper Perennial.

Murray, D., Sytsma, J., & Livengood, J. (2012). God knows (but does God believe?). *Philosophical Studies*, 166(1), 83–107.

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

Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86(2), 279–294.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84(3), 231.

Pinillos, N. Á. (2012). Knowledge, experiments and practical interests. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions*. Oxford: Oxford University Press.

Reed, B. (2013). Practical matters do not affect whether you know. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary debates in epistemology* (2nd ed., pp. 95–106). Malden, Mass.: Wiley-Blackwell.

Rose, D., & Nichols, S. (2013). The lesson of bypassing. *Review of Philosophy and Psychology*, 4, 599–619.

Rose, D., Livengood, J., Sytsma, J., & Machery, E. (2012). Deep trouble for the deep self. *Philosophical Psychology*, 25(5), 629–646.

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

Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.

Turri, J. (2016a). How to do better: toward normalizing experimentation in epistemology. In J.

Nado (Ed.), *Advances in experimental philosophy and philosophical methodology* (pp. 35–51). London: Bloomsbury Academic.

Turri, J. (2016b). Knowledge judgments in "Gettier" cases. In J. Sytsma & W. Buckwalter (Eds.), *A companion to experimental philosophy* (pp. 337–348). Wiley-Blackwell.

Turri, J. (2016c). *Knowledge and the norm of assertion: an essay in philosophical science*. Cambridge: Open Book Publishers.

Turri, J. (2016d). Vision, knowledge, and assertion. *Consciousness and Cognition*, *41*(C), 41–49. http://doi.org/10.1016/j.concog.2016.01.004

Turri, J., & Buckwalter, W. (in press). Descartes's schism, Locke's reunion: completing the pragmatic turn in epistemology. American Philosophical Quarterly.

Turri, J., Friedman, O., & Keefner, A. (in press). Knowledge central: a central role for knowledge attributions in social evaluations. *Quarterly Journal of Experimental Psychology*.

Weatherson, B. (2005). Can se do without pragmatic encroachment? *Philosophical Perspectives*, 19(1), 417–443.

Wilson, T. D. (2002). *Strangers to ourselves: discovering the adaptive unconscious*. Cambridge, Mass.: Harvard University Press.