



Error possibility, contextualism, and bias

Wesley Buckwalter¹ 

Received: 13 June 2018 / Accepted: 19 April 2019
© The Author(s) 2019

Abstract

A central theoretical motivation for epistemic contextualism is that it can explain something that invariantism cannot. Specifically, contextualism claims that judgments about “knowledge” are sensitive to the salience of error possibilities and that this is explained by the fact that salience shifts the evidential standard required to truthfully say someone “knows” something when those possibilities are made salient. This paper presents evidence that undermines this theoretical motivation for epistemic contextualism. Specifically, it demonstrates that while error salience does sometimes impact “knowledge” judgments as contextualism predicts, it does so in ways that are consistent with invariantism and does not require positing any additional contextualist semantics to explain. These results advance our understanding of the pathways by which error possibility affects “knowledge” judgments, answer a major challenge to invariantism, and suggest several methodological improvements for the study of knowledge attribution.

Keywords Knowledge · Evidence · Truth · Error salience · Contextualism · Mental state representation

1 Introduction

Chances are that if you are reading a paper on epistemic contextualism, you will encounter pairs of cases somewhere within the first few paragraphs. The purpose of these cases is to invite intuitive verdicts about the uses of “knowledge” sentences. These are taken throughout the literature to illustrate and theoretically motivate the view that “knowledge” is a context sensitive expression (Cohen 2002; DeRose 2009). They were instrumental in the initial publication and dissemination of the theory and have since been offered as a central theoretical motivation for contextualism over rival invariantist theories in epistemology (DeRose 2009, 1992). What are these intuitive

✉ Wesley Buckwalter
wesleybuckwalter@gmail.com

¹ Department of Philosophy, University of Manchester, Manchester, UK

verdicts about “knowledge” sentences and how do they uniquely motivate contextual semantics over invariantist semantics about the verb “knows”?

One central class of intuitive verdicts involve the effect that the salience of error possibility across pairs of cases has on judgments about the truth of “knowledge” sentences (Nagel 2010, 2012b; Cohen 2002; DeRose 1992, 2009; Lewis 1996; Gerken 2013; Hawthorne 2004; Williamson 2005; Cohen 1999). In a low error salience version of the case, the possibility that a subject has made an error is not made salient. In a high error salience version of the case, the situation is said to be exactly the same, except for the fact that a possibility of error has been made salient. Many philosophers have argued that we are likely to say that knowledge affirming sentences are true in the first case and that knowledge denying sentences are true in the second case. One famous example, for instance, involves discerning the color of a table in a furniture store. It is often claimed that we are less likely to judge that a subject “knows” that the table they are looking at is red when it is mentioned that there is a small chance that it is a white table illuminated by a light to merely appear red, than when no one raises that skeptical possibility (Cohen 2002; Nagel 2010). And in fact, such claims have inspired a wave of research in experimental cognitive science confirming that salience can impact “knowledge” judgments in this way, not only in professional philosophical judgment, but also in ordinary patterns of knowledge representation in social cognition (Alexander et al. 2014; Buckwalter and Schaffer 2015; Gerken and Beebe 2016; Nagel et al. 2013; Powell et al. 2015; Schaffer and Knobe 2012; for a partial review, see Buckwalter 2017).

Contextualists have argued that this effect of error salience on “knowledge” judgments uniquely motivates the view that “knows” is context sensitive (Lewis 1996; DeRose 2009; Cohen 2013). According to contextualism, the evidential standards required to truthfully say a person “knows” something are set by and shift according to the contexts in which that statement is made. If only the salience of error, and no other factors thought to be required for “knowing” by invariantists shift between pairs of salience cases, this line of thinking goes, then the variation in “knowledge” judgments must be explained by shifts in the standards for “knowing”. Hence, the theory can explain “knowledge” judgments when invariantism purportedly cannot because the salience of error possibility raises the evidential standard required to truthfully say someone knows something. It is this shift in the evidential standard that explains shifts in “knowledge” judgments across cases and provides the motivation to posit contextual semantics to explain.

In response, invariantists have proposed various explanations of why “knowledge” judgments differ in salience cases without positing shifting contextual standards. Many theories take the form of error theoretic or debunking style arguments, which employ research from the judgment and decision-making literature to hypothesize how professional philosophical intuition and ordinary judgment could be widely mistaken.¹ For example, one hypothesis is that attributions are the result of the “availability heuristic” (Tversky and Kahneman 1973), which causes us to misrepresent the odds an error has been made in high salience conditions (Williamson 2005; Hawthorne 2004). Alter-

¹ For other approaches, see skeptical invariantist responses (Unger 1975), or responses in terms of speech act theory or the pragmatics of assertion (Turri 2014; Brown 2005; Rysiew 2012).

natively, others hypothesize that the effect is explained by “epistemic egocentrism”, which causes us to misrepresent the mental states of subjects who possess less information than we do and subsequently deny knowledge to them (Nagel 2010, 2012b). A third hypothesis is that the effect is a “focal bias”, in which salient alternatives are mistakenly processed as epistemically relevant due to which alternatives are psychologically emphasized (Gerken 2013, 2017; Stanovich 2009).

Fundamental to this dialectic however is one crucial assumption made about how salience impacts “knowledge” judgments upon which the central motivation for positing contextualism over invariantism, and subsequently, the entire debate that occurred as a result of explaining error salience effects absolutely depends. The assumption involves the construction of salience cases and the pathway by which salience affects “knowledge” judgments. Clearly, error possibilities might affect judgments in all kinds of ways with no theoretical significance with respect to invariantism or contextualism. For example, one way to make the possibility of error salient is to suggest that a subject was sleep deprived or hallucinating when forming a belief. Alternatively, perhaps the strongest way to raise the possibility of error is to make this possibility actual, by saying that a belief was in fact made in error. These things would have no theoretical significance for the present debate however because they would involve manipulating not only salience, but invariantist factors as well, such as the amount of evidence or truth of the belief in question. Instead, the unique motivation for contextualism requires that salience cases hold all other variables equal and that salience affects “knowledge” judgments *directly* and *independently* of invariantist requirements. In other words, the motivation is predicated on the claim that variance in “knowledge” judgments cannot simply be explained by changes in other judgments across contexts, such as truth or evidence. If shifts in invariantist factors can explain the effect of salience on “knowledge” judgments than there is no need to posit shifting standards to explain that effect.

For this reason, the salience literature is full of assurances that salience case pairs are matched all-else-equal. We are told, for example, “two subjects of a knowledge ascription are on a par with regard to all truth-relevant parameters” (Gerken 2013). Others observe that “these cases do not necessarily describe different subjects, or subjects in different conditions,” and that the shift to high salience “simply adds some detail to the description” which involves “drawing attention to a fact we could have recognized from the outset as a feature of both cases” (Nagel 2010). However, these assurances concerning how cases are supposedly constructed do not tell us the psychological pathway of salience effects when they are processed or if they are actually held all-else-equal. The pathway matters because if the processing does not reflect what the cases are assumed to be like, and “knowledge” judgments fluctuate as a result of changes in other invariantist factors, then salience provides no motivation to posit contextualist semantics over existing invariantist accounts.

This paper evaluates whether this is the case. Specifically, the following research tests whether the effect that salience manipulations have on “knowledge” judgments in well-known cases developed by contextualists actually require contextualist over invariantist semantics to explain. The central research question is whether salience effects occur all-else-equal, or whether these effects on “knowledge” judgments are indirect and dependent on judgments about other variables that are taken by invari-

artists to be required for “knowing”. To evaluate this research question, classic contextualist case pairs are tested along with a series of probes designed to test for shifts in invariantist factors. Three truth-conducive invariantist factors were chosen for inclusion, including: the truth of a belief, whether a belief is held, and the amount of evidence for that belief. Additionally, one non-truth-conducive invariantist factor was also included due to the important role that subject-sensitive invariantist and pragmatist theories have played in recent debates about contextualism: actionability, or how one should act in a certain situation (Turri et al. 2016; Turri and Buckwalter 2017). If these factors are held fixed and “knowledge” judgments fluctuate as a result of manipulating error salience, then it provides evidence that uniquely motivates contextualism over rivals. If judgments about these things along with “knowledge” judgments vary, then it undermines the central motivation for contextualism. It undermines the central motivation for contextualism because even if forms of contextualism predict such findings, the findings are also predicted by rival theories even if contextualism were false. Depending on how those judgments are affected, it may also question extant invariantist error-theoretic accounts involving various psychological biases offered to explain those patterns.

This research is motivated by three prior advances in epistemology and cognitive science suggesting that cases are not held all-else-equal. First, several philosophers have argued that truth-conducive variables, such as belief or evidence, can explain the impact of non-truth-conducive variables, such as practical interests or stakes, on “knowledge” judgments (Weatherson 2005; Bach 2005; Nagel 2008; see also for discussion Turri and Buckwalter 2017). This raises the possibility that they could also mediate the impact of salience on “knowledge” judgments. Second, it has been demonstrated that both judgments about quality of evidence and propriety of beliefs are truth-sensitive (Turri 2016). If a belief actually being mistaken impacts judgments about belief or evidence one has for that belief, then it is reasonable to suppose that the possibility a belief is mistaken might also impact judgments about quality and propriety of belief or evidence. Third, prior research has demonstrated that epistemic variables shift in other classic contextualist thought experiments, such as bank cases (DeRose 1992) and airport cases (Cohen 1999) when stakes and error are manipulated (Turri 2017). This questions the motivation for contextualism offered by these cases, though prior research has not studied the pathway of salience effects specifically. In what follows, three behavioral experiments are conducted to directly test whether prior salience cases presented in the contextualist literature are direct or indirect, and thus motivate contextualism.

2 Experiment 1: red light cases

The following experiment tests whether the effect of salience on “knowledge” judgments is both indirect and dependent on invariantist factors. In this experiment, ninety-nine participants were tested, aged 19–69 years, mean age = 36 years; 51 female; 91% reporting English as a native language. Participants were recruited and tested online using Amazon Mechanical Turk and Qualtrics, and compensated \$0.35 for approximately two minutes of their time. Repeat participation was prevented within

and across experiments. Participants were randomly assigned to one of two salience conditions from the famous red-light cases introduced above. The text of the cases is verbatim from Nagel (2010) and prior stimuli administered experimentally by Alexander et al. (2014). The low salience case read:

(Low Light) John A. Doe is in a furniture store. He is looking at a bright red table under normal lighting conditions. He believes the table is red.

The high salience condition read:

(High Light) John B. Doe is in a furniture store. He is looking at a bright red table under normal lighting conditions. He believes the table is red. However, a white table under red lighting conditions would look exactly the same to him, and he has not checked whether the lighting is normal, or whether there might be a red spotlight shining on the table.

The story remained at the top of the screen while participants evaluated five statements presented in random order involving truth, belief, evidence, knowledge, and actionability, involving an action that pertained specifically to the further collection of evidence:

It's true that the table is red. (truth)

John believes that the table is red. (belief)

John has good evidence that the table is red. (evidence)

John should ask a sales associate if the table is red. (actionability)

John knows that the table is red. (knows)

Responses were collected on a standard seven-item Likert agreement scale, anchored “strongly disagree” to “strongly agree”, left-to-right across the participant’s screen.²

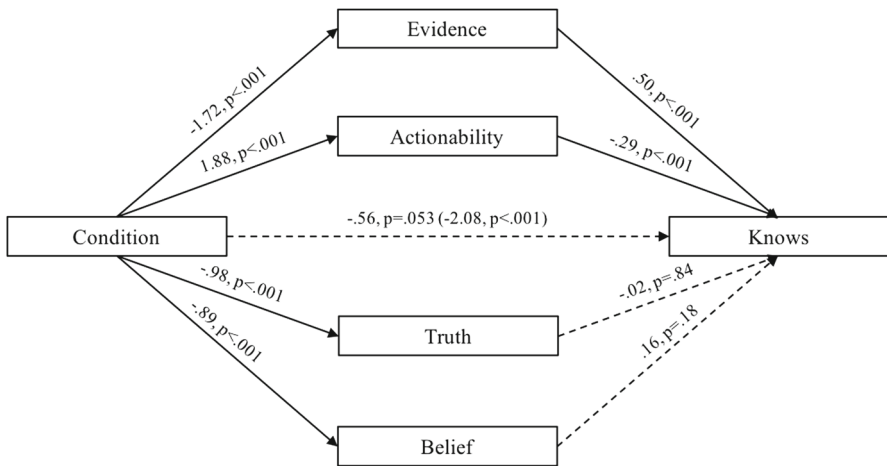
Results indicate that assignment to condition significantly affected responses to all statements. Participants in the high salience condition were less likely to agree that it is true, that there is good evidence, and that John believes and knows that the table is red, and were more likely to agree that John should collect more evidence that the table is red (see Table 1). According to conventional effect size guidelines (Ellis 2010), the effect on knowledge statements was large ($d = 1.47$). Mean agreement with knowledge statements was significantly greater than chance (neutral midpoint = 4) in low table, $t(48) = 11.20$, $p < .001$, $MD = 1.96 [1.61, 2.31]$, $d = 1.60$, but was not distinguishable from chance rates in high table, $t(49) = -.54$, $p = .595$, $MD = -0.12 [-.57, .33]$, $d = -.08$.

To understand the relationship between these variables, a bootstrap mediation analysis was conducted (Hayes 2013). Mediation analysis is clearly relevant to test the proposal that salience affects “knowledge” judgments independently of shifting invariantist factors. Mediation analysis is a technique for clarifying how a predictor influences an outcome. This analysis estimates how much of the predictor’s influence on the outcome is direct and how much influence is indirect. Indirect influence is

² Given that agreement bias has been documented by prior researchers in contextualist cases (Turri 2017), all experiments reported in the paper use a uniform phrasing of the knowledge statement across low and high salience conditions.

Table 1 Experiment 1: descriptive statistics along with analysis of difference for all measures in low and high light conditions

Measure	Low M (SD)	High M (SD)	t	df	p	MD	95% CI	d
Truth	5.98 (1.01)	5.00 (1.54)	3.74	97	<.001	0.98	0.46, 1.50	0.75
Belief	6.59 (0.64)	5.70 (1.33)	4.24	97	<.001	0.89	0.47, 1.31	0.85
Evidence	6.12 (0.88)	4.40 (1.47)	7.05	97	<.001	1.72	1.24, 2.21	1.42
Actionability	3.06 (1.77)	4.94 (1.82)	-5.20	97	<.001	-1.88	-2.60, -1.16	-1.05
Knows	5.96 (1.22)	3.88 (1.59)	7.29	97	<.001	2.08	1.51, 2.65	1.47

**Fig. 1** Experiment 1: mediation of the relationship between assignment to salience condition and knowledge in red light cases, $R^2 = .64$, $p < .001$

mediated by other variables, whereas direct influence is not. As explained above, a unique motivation for contextualism requires not only that there is a salience effect on “knowledge” judgments, but also that this effect is direct. In the present bootstrap mediation model, knows was treated as the outcome variable, assignment to salience condition as the independent variable (coded: low light = 0, high light = 1), and all other dependent variables as potential mediators. The analysis showed that evidence and actionability were significant mediators, indirect effects = $-.86$ [$-1.46, -.27$], $-.54$ [$-1.10, -.19$], that belief and truth were not significant mediators, and that the model fully accounted for the effect of salience condition on responses to knowledge statements (see Fig. 1).

These results demonstrate four things. First, mention of error possibility had a large effect on responses to knowledge statements, replicating findings from prior researchers (Alexander et al. 2014). Instead of being held fixed across contexts however, salience impacted all judgments measured in the experiment, including what is thought to be true, believed, the evidence quality, and actionability. Because judgments about invariantist requirements for knowledge shift in these pairs of cases, there is no theoretical pressure to adopt a semantic account in which the standards for “knowl-

edge” shift between them. Second, though the effect of salience on knowledge was large, the central tendency was not for participants to disagree with knowledge statements in the high salience condition. Participants in this case were ambivalent. Third, judgments about evidence and actionability were the best predictors of responses to knowledge statements. Fourth, shifts in judgments about invariantist factors can fully explain the effect of salience on responses to knowledge statements.

3 Experiment 2: diamond cases

To ensure that these results are generalizable, a conceptual replication was performed using the text of another thought experiment in the salience literature. Two hundred and one new participants were tested using the same procedures described above, aged 20–71 years, mean age = 33 years; 89 female; 94% reporting English as a native language. Participants were randomly assigned to one of two salience conditions with text taken verbatim from Nagel et al. (2013). The low salience case read:

(Low Diamond) Emma is shopping for jewelry. She goes into a nice-looking store, and spends some time looking at various different displays. She tells the salesperson that she is looking for a simple diamond necklace with a classic design. She always likes to try things on before she makes up her mind about them, and she asks the salesperson to show her quite a wide variety of different items, which he brings out for her one tray at a time. Emma selects a diamond necklace from a tray marked “Diamond Earrings and Pendants”. “What a lovely diamond!” she says as she tries it on.

The high salience case read:

(High Diamond) Emma is shopping for jewelry. She goes into a nice-looking store, and spends some time looking at several different displays. She tells the salesperson that she is looking for a diamond necklace with a classic design. She always likes to try things on before she makes up her mind about them, and the salesperson shows her several options. Emma selects a diamond pendant from a tray marked “Diamond Earrings and Pendants”. “What a lovely diamond!” she says as she tries it on. Emma could not tell the difference between a real diamond and a cubic zirconium fake just by looking or touching.

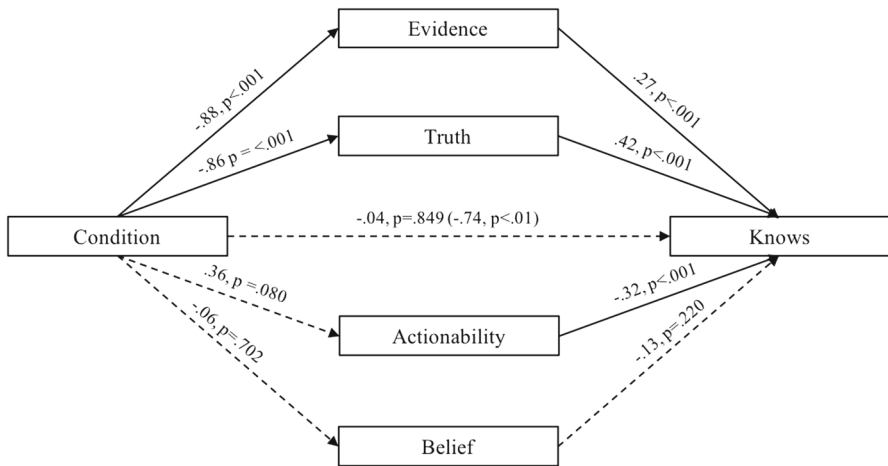
After reading one of these two stories, participants evaluated a similar set of statements to those above, including:

- It’s true that the stone in the neckless is a diamond. (truth)
- Emma thinks that the stone in the necklace is a diamond. (belief)
- Emma has good evidence that the stone in the necklace is a diamond. (evidence)
- Emma should double check that the stone in the necklace is a diamond. (actionability)
- Emma knows that the stone in the necklace is a diamond. (knows)

Results indicate that salience significantly impacted truth, evidence, and knowledge judgments. More specifically, participants in high diamond were less likely to agree

Table 2 Experiment 2: descriptive statistics along with analysis of difference for all measures in low and high diamond conditions

Measure	Low M (SD)	High M (SD)	t	df	p	MD	95% CI	d
Truth	5.08 (1.17)	4.22 (1.58)	4.38	199	<.001	0.86	0.47, 1.24	0.62
Belief	6.17 (1.03)	6.11 (1.03)	0.38	199	.702	0.06	- 0.23, 0.34	0.06
Evidence	5.15 (1.45)	4.26 (1.60)	4.12	199	<.001	0.88	0.46, 1.31	0.58
Actionability	5.28 (1.54)	5.65 (1.37)	- 1.76	199	.080	- 0.36	- 0.77, 0.04	- 0.25
Knows	4.44 (1.81)	3.70 (1.66)	3.04	199	.003	0.74	0.26, 1.23	0.43

**Fig. 2** Experiment 2. Mediation of the relationship between assignment to salience condition and knowledge in diamond cases, $R^2 = .38, p < .001$

that the stone in the necklace is a diamond, that Emma has good evidence, and that Emma knows that the stone is a diamond (see Table 2). There was a small effect on responses to knowledge statements ($d = .43$). Mean agreement with knowledge statements were significantly greater than chance (neutral midpoint = 4) in low diamond, $t(101) = 2.47, p < .05, MD = 0.44 [0.09, .80], d = .24$, but was not distinguishable from chance rates in high diamond, $t(98) = -1.81, p = .07, MD = -0.30 [-.63, .03], d = .18$.

A bootstrap mediation analysis was again conducted with knows as the outcome variable, assignment to salience condition as the independent variable (coded: low diamond = 0, high diamond = 1), and all other dependent variables as potential mediators. The analysis showed that evidence and truth were significant mediators, indirect effects = $-.24 [-.46, -.08], -.36 [-.62, -.18]$. Belief and actionability were not significant mediators. This model fully accounted for the effect of condition on responses to knowledge statements (see Fig. 2).

These results demonstrate five things. First, assignment to condition affected responses to knowledge sentences, replicating prior research (Nagel et al. 2013). Second, invariantist factors were not held fixed between cases. In addition to affecting

responses to knowledge statements, salience also impacted judgments concerning what was thought to be true in the case itself and strength of the evidence. Third, participants again agreed with knowledge statements in the low salience condition but were ambivalent in the high salience condition. The central tendency was not to disagree with knowledge statements. Fourth, judgments about evidence and truth were the best predictors of knowledge. Fifth, invariantist factors fully explain the effect of condition on responses to knowledge statements.

4 Experiment 3: clock cases

To ensure further generalizability, a third experiment was performed using other narrative materials from the salience literature. Two hundred and three new participants were tested using the same procedures described above (aged 20–89 years, mean age = 36 years; 108 female; 93% reporting English as a native language). Participants were randomly assigned to one of two conditions taken verbatim from Nagel (2012a). The low salience case read:

(Low Clock) Wanda is out for a weekend afternoon walk. She lives in a large new condominium tower downtown, and her suite is fairly small and does not have any windows that open, so she really likes to get out for some fresh air. Passing near the train station, Wanda wonders what time it is. She glances up at the clock on the train station wall and sees that it says 4:15 pm. It is in fact 4:15 pm at that moment.

The high salience case read:

(High Clock) Wanda is out for a weekend afternoon walk near the train station and wonders what time it is. She glances up at the clock on the train station wall and sees that it says 4:15 pm. It is in fact 4:15 pm at that moment. The station clock is in fact working, but it has no second hand, and Wanda only looks at it for a moment, so she would not be able to tell if the clock were stopped.

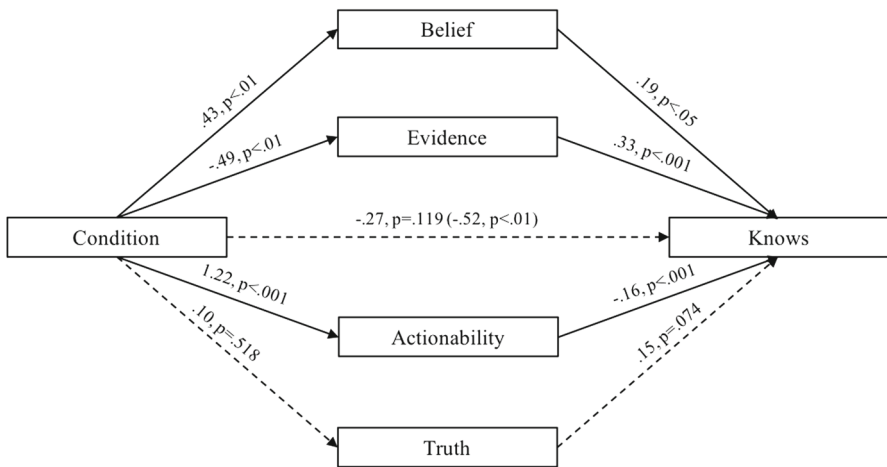
After reading one of these two stories, participants evaluated statements using a matrix table. The top of the table was prefaced by the statement, “At the moment Wanda sees the clock:” followed by the statements:

- It’s true that that it is 4:15 pm. (truth)
- Wanda thinks that it is 4:15 pm. (belief)
- Wanda has good evidence that it is 4:15 pm. (evidence)
- Wanda should double check whether the clock is working. (actionability)
- Wanda knows that it is 4:15 pm. (knows)

Results indicate that assignment to condition affected responses to belief, evidence, actionability, and knowledge statements. More specifically, participants in the high salience condition were less likely to agree that Wanda believes, has good evidence, and knows the time, and more likely to agree she should collect more evidence (see Table 3). The effect on knowledge statements was small ($d = .41$). Mean agreement with knowledge statements was significantly greater than chance (neutral midpoint =

Table 3 Experiment 3: descriptive statistics along with analysis of difference for all measures in low and high clock conditions

Measure	Low M (SD)	High M (SD)	t	df	p	MD	95% CI	d
Truth	6.29 (1.08)	6.38 (1.02)	-0.65	201	.518	-0.10	-0.39, 0.20	-0.09
Belief	5.98 (1.36)	6.41 (0.77)	-2.78	201	<.01	-0.43	-0.74, -0.13	-0.39
Evidence	6.14 (1.17)	5.66 (1.37)	2.73	201	<.01	0.49	0.14, 0.84	0.38
Actionability	3.02 (1.61)	4.24 (1.84)	-5.06	201	<.001	-1.22	-1.70, -0.75	-0.71
Knows	5.87 (1.23)	5.34 (1.38)	2.85	201	<.01	0.52	0.16, 0.88	0.41

**Fig. 3** Experiment 3. Mediation of the relationship between assignment to salience condition and knowledge in clock cases, $R^2 = .34$, $p < .001$

4) in both low clock, $t(103) = 15.45$, $p < .001$, $MD = 1.87$ [1.63, 2.10], $d = 1.52$, and high clock, $t(98) = 9.69$, $p < .001$, $MD = 1.34$ [1.07, 1.62], $d = .97$, indicating strong agreement in both cases.

A bootstrap mediation analysis was conducted with knows as the outcome variable, assignment to condition as the independent variable (coded: low clock = 0, high clock = 1), and all other dependent variables as potential mediators. The analysis showed that belief, indirect effect = .08 [.01, .22], evidence, indirect effect = -.16 [-.36, -.04], and actionability, indirect effect = -.19 [-.41, -.07], were significant mediators. Truth was not a significant mediator. This model fully accounted for the effect of salience on responses to knowledge statements (see Fig. 3).

These results demonstrate that invariantist factors were not held fixed when manipulating salience. Though a salience effect was detected on responses to knowledge statements, the central tendency was agreement in low and high salience conditions. Lastly, replicating the findings above, a mediation model including all dependent variables fully explains the effect of salience on responses to knowledge statements.

5 Conclusion

Philosophers and cognitive scientists disagree about the nature of salience effects and their theoretical significance in epistemology. This research further improves our understanding of the effect and how it should be interpreted. Salience was found to have an effect on “knowledge” judgments, ranging from small to large in size depending on the case. This never resulted in a central tendency that flipped reactions to knowledge sentences from agreement to disagreement, but did sometimes reduce clear agreement with knowledge sentences to indecision. In other instances, the effect was found despite agreement with knowledge sentences in both low and high salience conditions. Perhaps most importantly for understanding their theoretical significance however, the present findings demonstrate two things. First, salience manipulations not only impacted “knowledge” judgments, but several invariantist factors, including what is considered true, believed, good evidence, and actionable. Manipulations did more than add detail to the description that could have been recognized about both contexts, but in effect, precipitated different invariantist contexts. Second, not only did these factors shift, but what is more, those shifts in invariantist factors even explained variance in “knowledge” judgments as a result of error salience.

These findings suggest that the salience effects on “knowledge” judgments detected are indirect and dependent on truth-conducive and non-truth conducive factors also predicted by invariantism. This undermines the motivation that salience effects have been taken to offer contextualism over invariantism. It undermines this by demonstrating that while it may or may not be true that salience effects can be explained by contextualism, they are also just as well explainable by rival theories even if contextualism is false. If this is true, then there is no need to posit shifting standards because variation in “knowledge” judgements can be explained by conventionally recognized invariantist factors. Specifically, these patterns are explained by the fact that “knows” is invariant, that it requires things like truth, belief, and evidence, and that judgments about those things differ across contexts. Thus, if salience effects can also be explained by invariantism, they cannot motivate contextualism over invariantism.

Though the present research demonstrates that salience effects are indirect in some of the best thought experiments provided in its favor, this does not rule out the possibility of discovering a direct effect. To move this line of research forward productively in that direction, researchers should do four things. First, devise more carefully worded cases that constitute minimally matched pairs and that clearly specify shared invariantist components. Given the results above, researchers should be especially careful that error manipulations do not inadvertently undermine truth or the quality of evidence. Second, demonstrate that “knowledge” judgment shifts as predicted. Third, verify and control for the effect of shifting invariantist factors, as well as likely moderators, such as wording effects or task demands. Fourth, provide evidence that “knowledge” judgments cannot be adequately explained by shifts in those variables. Until such evidence for direct effects is found, researchers cannot assume salience effects motivate contextualism over invariantism.

The support the findings offer existing hypotheses of psychological bias in the salience literature is mixed. To the extent that error theoretic accounts do not recognize that belief, truth, and evidence judgments each fluctuate, they do not fully explain

how “knowledge” judgments are affected by error salience. At the same time, however, future research might potentially deploy extant accounts at the level of each of those judgments. For instance, perhaps it is not “knowledge” judgment, but rather, assessments of belief, truth, and evidence that are each explained by psychological bias when salience is manipulated, and it is this suite of biased judgments that then impact responses to knowledge statements as epistemologists have traditionally thought belief, truth, and evidence should do that. Further research might explore whether the various biases that have been proposed can be applied at the level of each of those judgments.

Though this is an open theoretical possibility, it is not obvious what the present motivation is for doing so. For one, there is little theoretical pressure to discount these judgments, since the resulting asymmetry in “knowledge” judgments they explain does not motivate contextualism over invariantism. For another, a singular psychological bias is unlikely to explain variance in many of the judgments observed across the broad range of salience cases that predict “knowledge” judgements. For example, epistemic egocentrism is primarily a hypothesis involving misrepresenting the mental states of others, not what is true. Relatedly, the availability heuristic account is a view involving the probability of outcomes, not what an agent believes. Nonetheless, both truth and belief judgments shift in various experimental contexts. Further research might explore whether prior accounts could overcome these difficulties.

Independently of the specific accounts of bias that have been proposed by invariantists, a separate question is whether epistemic judgments in the cases above are mistaken more generally, in the sense that they violate theoretical conventions in traditional epistemology. Put another way, is it a mistake to regard them as different epistemic contexts? Though this broader question is beyond the full scope of this paper, when examining the wordings in detail, it is not initially obvious it is a mistake. For one, the wordings allow for charitable disagreement about how cases are being interpreted. For another, mediation analysis displays a series of fairly orthodox epistemic patterns. When truth mediates the effect of salience on responses to knowledge statements, in the diamond case, for instance, it is not made as clear as it could be what is true. It might be reasonable to doubt the stone is a diamond as presented. When belief mediates the effect of salience on responses to knowledge statements, such as in the clock case, it is never explicitly specified what is believed. And when evidence plays a mediating role, as in red light cases, the evidence is somewhat unclear. Perhaps it is reasonable to doubt that the presence of table-altering spotlights couldn’t have been detected. Or perhaps it is reasonable to be suspicious if someone “has not checked” things we typically assume they know, such as whether there are lights shining on things in stores. Relatedly, it might be thought, an agent looking at and seeing an object while not seeing something else about that very same object gives new information that should result in ambivalence about the reliability of that perceptual relation and resulting knowledge sentences. Though further research is necessary to explore these normative concerns about what judgments about epistemic contexts should be like, these contingencies raise the possibility that invariantist explanations may not require positing as much error among ordinary and professional “knowledge” judgments as some researchers have supposed (Williamson 2005; Gerken 2017). Instead, these contingencies raise the following hypothesis that seems just as likely to explain the present effects as

contextualism or hypotheses involving psychological bias do. Rather than contextualist semantics or systematic cognitive biases, the effects on “knowledge” judgments observed may result from a cluster of experimenter and wording effects introduced in the construction of thought experiments that tacitly underspecify or obfuscate key invariantist requirements for knowledge.

Lastly, this research suggests several opportunities to improve research on knowledge attribution. The first is that when relying on psychological judgments to support a theory, such as those involving the truth of knowledge sentences or knowledge attribution, the theoretical significance of such judgments sometimes depends on the pathway that generates them. Second, utilizing the tools of experimental cognitive science, such as mediation analysis, can help researchers to understand this philosophical significance. This is especially true when philosophical arguments depend on controlling for multiple variables and precise manipulations. Third, drawing on extant psychological literature about related but distinct areas of mental processing can constitute a promising start to inquiry aimed at understanding judgments in different domains. However, this can also potentially lead research astray if further theorizing occurs as a substitute for a direct test of the initial hypothesis in the target domain. Finally, rather than looking for reasons to accept or dismiss judgments as a means to advance prior theoretical commitments, researchers might suspend judgment about what these behaviors support until the mechanisms are better understood to minimize the risk of theorizing beyond our understanding of knowledge sentences.

Acknowledgements For helpful feedback, discussion, and comments on previous drafts, I thank Joshua Alexander, Marc Alspector-Kelly, Carolyn Buckwalter, Mikkel Gerken, Edouard Machery, and John Turri.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Alexander, J., Gonnerman, C., & Waterman, J. (2014). Salience and epistemic egocentrism: an empirical study. In J. Beebe (Ed.), *Advances in experimental epistemology* (pp. 97–118). London: Bloomsbury.
- 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.
- Brown, J. (2005). Comparing contextualism and invariantism on the correctness of contextualist intuitions. *Grazer Philosophische Studien*, 69(1), 71–100. <https://doi.org/10.1163/18756735-069001005>.
- Buckwalter, W. (2017). Epistemic contextualism and linguistic behavior. In J. J. Ichikawa (Ed.), *The routledge handbook of epistemic contextualism* (pp. 44–56). New York: Routledge.
- Buckwalter, W., & Schaffer, J. (2015). Knowledge, stakes, and mistakes. *Nous*, 49(2), 201–234.
- Cohen, S. (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, 13, 57–89.
- Cohen, S. (2002). Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, 65(2), 309–329. <https://doi.org/10.1111/j.1933-1592.2002.tb00204.x>.
- Cohen, S. (2013). Contextualism Defended. In M. Steup, J. Turri, & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (2nd ed., pp. 69–74). Malden, MA: Wiley-Blackwell.
- DeRose, K. (1992). Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4), 913–929.
- DeRose, K. (2009). *The case for contextualism*. Oxford: Oxford University Press.

- Ellis, B. D. (2010). *The essential guide to effect sizes: Statistical power, meta-analysis, and the interpretation of research results*. Cambridge: Cambridge University Press.
- Gerken, M. (2013). Epistemic focal bias. *Australasian Journal of Philosophy*, 91(1), 41–61.
- Gerken, M. (2017). *On folk epistemology: How we think and talk about knowledge*. New York, NY: Oxford University Press.
- Gerken, M., & Beebe, J. R. (2016). Knowledge in and out of Contrast. *Nous*, 50(1), 133–164. <https://doi.org/10.1111/nous.12064>.
- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Hayes, A. F. (2013). *Introduction to mediation, moderation, and conditional process analysis: A regression-based approach*. New York: Guilford Press.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–567.
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86(2), 279–294. <https://doi.org/10.1080/00048400801886397>.
- Nagel, J. (2010). Knowledge ascriptions and the psychological consequences of thinking about error. *The Philosophical Quarterly*, 60(239), 286–306. <https://doi.org/10.1111/j.1467-9213.2009.624.x>.
- Nagel, J. (2012a). Intuitions and experiments: A defense of the case method in epistemology. *Philosophy and Phenomenological Research*, 85(3), 495–527.
- Nagel, J. (2012b). Mindreading in gettier cases and skeptical pressure cases. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions* (pp. 171–191). Oxford: Oxford University Press.
- Nagel, J., San Juan, V., & Mar, R. A. (2013). Lay denial of knowledge for justified true beliefs. *Cognition*, 129, 652–661.
- Powell, D., Horne, Z., Pinillos, Á., & Holyoak, K. (2015). A Bayesian framework for knowledge attribution: Evidence from semantic integration. *Cognition*, 139, 92–104.
- Rysiew, P. (2012). Epistemic scorekeeping. In J. Brown & M. Gerken (Eds.), *Knowledge ascriptions*. Oxford: Oxford University Press.
- Schaffer, J., & Knobe, J. (2012). Contrastive knowledge surveyed. *Nous*, 46(4), 675–708.
- Stanovich, K. E. (2009). Distinguishing the reflective, algorithmic, and autonomous minds: Is it time for a tri-process theory. In K. Frankish & J. S. B. T. Evans (Eds.), *In two minds: Dual processes and beyond* (pp. 55–88). Oxford: Oxford University Press.
- Turri, J. (2014). Linguistic Intuitions in Context: A Defense of Non-Skeptical Pure Invariantism. In A. Booth & D. Rowbottom (Eds.), *Intuitions* (pp. 165–184). Oxford: Oxford University Press.
- Turri, J. (2016). The radicalism of truth-insensitive epistemology: Truth's profound effect on the evaluation of belief. *Philosophy and Phenomenological Research*, 93(2), 348–367. <https://doi.org/10.1111/phpr.12218>.
- Turri, J. (2017). Epistemic contextualism: An idle hypothesis. *Australasian Journal of Philosophy*, 95(1), 141–156.
- Turri, J., & Buckwalter, W. (2017). Descartes's schism, Locke's reunion: Completing the pragmatic turn in epistemology. *American Philosophical Quarterly*, 54(1), 25–46.
- Turri, J., Buckwalter, W., & Rose, D. (2016). Actionability judgments cause knowledge judgments. *Thought*, 5(3), 212–222.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5(2), 207–232. [https://doi.org/10.1016/0010-0285\(73\)90033-9](https://doi.org/10.1016/0010-0285(73)90033-9).
- Unger, P. (1975). *Ignorance: A case for skepticism*. Oxford: Oxford University Press.
- Weatherson, B. (2005). Can we do without pragmatic encroachment? *Philosophical Perspectives*, 19(1), 417–443. <https://doi.org/10.1111/j.1520-8583.2005.00068.x>.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *The Philosophical Quarterly*, 55(219), 213–235. <https://doi.org/10.1111/j.0031-8094.2005.00396.x>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.