

Intentional Action Without Knowledge

Romy Vekony, Al Mele and David Rose

[forthcoming in *Synthese*]

In order to be doing something intentionally, must one know that one is doing it? Some philosophers have answered yes. Stuart Hampshire (1956: 95) contends that “if a man is doing something without knowing that he is doing it, then it must be true that he is not doing it intentionally.” Similarly, Christopher Olsen (1969: 324) claims that one “cannot not know” what one is doing intentionally. And Michael Gorr and Terence Horgan (1982: 255) claim that “P’s A-ing at t is intentional under the description ‘A-ing’ if and only if (i) this event is an act, and (ii) P knows, at t, of this act that it is an A-ing by him.” More recently, Eric Marcus (2019: 4) defends what he calls “The Knowledge Thesis” – “It is impossible for a person to do something intentionally without knowing that she is doing it” – against objections featuring self-deception. And John Schwenkler (2019: 34) attributes the following thesis to Elizabeth Anscombe: “A person does something intentionally only if she knows that she is doing it.” The quoted claims, straightforwardly interpreted, express or entail a thesis that we formulate as follows: (KT) An agent is doing something intentionally only if he knows that he is doing it.

A theorist who feels the pull of KT, may worry that requiring knowledge is overly bold. Perhaps awareness of what one is doing would suffice, even if this awareness sometimes falls short of knowledge. So, we test a more cautious version of KT, one that is entailed by KT but does not entail KT. Here it is: (KAT) An agent is doing something intentionally only if he knows that he is doing it or is aware that he is doing it. KAT does have an intuitive ring. But is it true? Here, using vignettes featuring skilled action and vignettes featuring habitual action, we provide evidence that, in various scenarios, a majority of non-specialists regard agents as intentionally doing things that the agents do

not know they are doing and are not aware of doing. Although we do not believe that the judgments of a majority of non-specialists settle whether KAT is true or false, we do believe that these judgments put pressure on proponents of KAT (and, of course, proponents of theses that entail KAT) to find a plausible way for our results to coexist peacefully with KAT's truth.

1. Setting the stage

The fact that there are competing theories of action individuation can complicate exposition. If by flipping the switch, Don turns on a light, illuminates a room, and unwittingly alerts a prowler, how many actions has he performed? Some philosophers (coarse-grained theorists; see Davidson 1963) say one action under four different descriptions. Others (fine-grained theorists; see Goldman 1970) say at least four different actions. Although lay folk are unaware of these competing theories, many of our readers will find them familiar. Fine-grained theorists among our readers should read KAT just as it is. Coarse-grained theorists should read it as follows: Under a description D, an agent is doing something intentionally only if, under D, he knows that he is doing it or is aware that he is doing it.

Consider the following from Donald Davidson's "Agency": "A man may even be doing something intentionally and not know that he is; so of course he can be doing it without knowing that he is. (A man may be making ten carbon copies as he writes, and this may be intentional; yet he may not know that he is; all he knows is that he is trying.)" (1980: 50). If we add that the man is not aware of making ten carbon copies, is this a successful counterexample to KAT? One can certainly imagine a philosopher resisting the claim that it is. How much control does the man have over whether he succeeds in making ten copies? Is his success too much a matter of luck for it to be true that he intentionally makes ten copies? Or is making ten carbon copies so easy for him that it would be wrong to say that he does not know that he is making ten copies? We did not wish to run Davidson's story by non-specialists; many potential participants simply are not familiar with carbon copies. But we did come up with some simple cases of our own.

There are some cases that one might think bear directly on KAT (see e.g., Lanteri 2009: 718-19 and Nadelhoffer 2004: 280-81). For instance, Lanteri (2009) uses a modified form of Knobe's (2003) original CEO cases but added that the CEO does not know what the environmental consequences are of the proposed plan:

Harm. An executive of a company goes to the chief executive officer and tells him 'We are thinking of launching a new product, which will increase profits. This product requires a new technique and we do not know what consequences this will have on the environment.' The CEO replies 'I do not care at all about the consequences. I just want to make profits. We shall launch the new product immediately.' The new product is launched, and as predicted, profits increase. However, the new technique turns out to be polluting and it harms the environment.

Help. [The same, except that the last sentence read 'Moreover, the new technique turns out to be ecological and it helps the environment.'.]

Participants were then asked the following two questions:

According to you, the CEO harmed/helped the environment...

(a) intentionally (b) not intentionally (c) neither

According to you, for his decision, the CEO should be...

(a) praised (b) blamed (c) neither (Lanteri 2009: 718)

The results showed that in *Harm*, 58% of respondents answered that the CEO intentionally harmed the environment, while in *Help*, only 7% thought that the CEO acted intentionally (Lanteri 2009: 719). While these results might appear to bear on KAT, it is not clear whether participants thought the CEO knew or was aware that starting the program would harm or help the environment. Since knowledge and awareness were not measured, we have no way of knowing whether people ascribed intentionality while at the same time denying that the CEO knew or was aware of the harmful

or helpful outcome. Aside from this, Lanteri is focused on both side-effect cases and on cases where the outcome is valenced. Proponents of KAT might view side-effect cases as not being central cases of intentional action and moreover might dismiss the relevance of these results on the basis of their arising as part of bias due to participants reactions to outcome valence (see e.g., Alicke and Rose 2010, Alicke and Rose 2012; Alicke, Rose and Bloom 2011). To assess KAT, our strategy will be to use non-side-effect cases that do not involve valenced outcomes. And we will ask participants whether the agent knew the outcome would occur, was aware that it would occur, and intentionally brought it about.

Our first experiment presents cases of highly skilled action in which the agent is precluded from access to sensory observation of successfully performing the action. We chose a case of highly skilled action so it was made clear to respondents that in the scenario, the successful performance of the action was far more likely due to the agent's skill rather than luck or causal deviance. The second experiment presents cases featuring habitual actions that the agent is unaware of performing. We chose to run the second study with a different action type to ensure that our results are generalizable beyond simply one action type. In both experiments, participants are asked to what extent they agree that the agent: (1) knows that he or she is performing the action, (2) is aware of performing the action, and (3) performed the action intentionally. If our results show that people are consistently willing to ascribe intentionality to an action while recognizing that the agent has no knowledge or awareness of performing that action, this poses a threat to KAT. But our concern is less about whether KAT is true or false than with how the folk come down on this. We seek to add to the growing body of knowledge about folk conceptions of intentional action. Our aim here is modest. It is to contribute to this growing body of knowledge by gathering evidence specifically about whether or not KAT is supported by lay responses to directly relevant cases. We do not explore alternatives to KAT here; doing so is a project for another occasion.

2. KAT empirically assessed

2.1. Study 1

Our first study focused on skilled action. Two hundred and fifty people, all adult residents of the United States, participated in this study. Participants were recruited from Amazon Mechanical Turk (<https://www.mturk.com>) and tested in Qualtrics (<https://www.qualtrics.com>).¹

Participants were randomly assigned to one of five scenarios. They first read a brief scenario and then responded to three test statements. The scenarios differed in many incidental details, from the agent's predicament to the action of interest. However, the scenario factor was not of independent theoretical interest and was included to support generalization of the results beyond the specific stimuli studied here (Baayen, Davidson, and Bates 2008; Clark 1973; Judd, Westfall, and Kenny 2012).

Here is one scenario (see Appendix for full set of scenarios):

[Basketball] Andy is a 92% free throw shooter. One evening, he is at the gym practicing his free throws. He lines up and takes the shot. But just as the ball leaves his hands, lightning strikes the building. The power goes out and it is pitch black. There is also a loud clap of thunder. Due to this, Andy could not see or even hear whether he made the shot. He is completely unaware of whether he sank the shot. But he did in fact sink the shot.

Beneath the scenario, on the same page, participants responded to three test statements (presented in random order):

When Andy was sinking the shot, he knew that he was sinking it. (Knowledge)

While Andy was sinking the shot, he was aware of sinking the shot. (Awareness)

Andy intentionally sank the shot. (Intentionality)

Responses to all test statements were collected on a standard 7-point Likert scale, 1 ("strongly disagree") - 7 ("strongly agree"), arranged left-to-right on the participant's screen.

¹ For both studies reported, only Mechanical Turk Workers who had a 95% or higher HIT approval rate were allowed to participate.

2.2. Results for Study 1

We conducted a Mixed ANOVA with Question (Knowledge, Awareness, Intentionality) as a within subject factor and Scenario (Basketball, Golf, Hockey, Pool, Air Hockey) as a between subject factor. There was a main effect of Question, $F(2, 490)=297.100, p<.001, \eta^2=.548$ and Scenario, $F(1, 245)=2.537, p<.041, \eta^2=.040$ but no Question by Scenario interaction, $F(8, 490)=1.541, p=.140, \eta^2=.025$.² Importantly, pairwise comparisons—using a Bonferonni correction of .016— indicated that Knowledge ($M=2.38, SD=1.74$) was not significantly different from Awareness ($M=2.20, SD=1.69$), $p=.036$ and Intentionality ($M=5.06, SD=1.96$), $p<.001$, and that Awareness was significantly different from Intentionality, $p<.001$.

One sample t-tests indicated that mean intentionality attribution was significantly above the midpoint in all scenarios except the Golf scenario while mean knowledge and awareness attributions were significantly below the midpoint in all scenarios (see Table 1).

Table 1
Experiment 1. One sample t-tests. Test value = 4

Measure	N	M	SD	df	t	p	d
<i>Basketball</i>							
Intentionality	50	5.10	2.00	49	3.88	<.001	0.55
Knowledge	50	2.68	1.83	49	-5.08	<.001	-0.72
Awareness	50	2.68	1.81	49	-5.15	<.001	-0.72
<i>Golf</i>							
Intentionality	50	4.40	2.23	49	1.26	.211	0.18
Knowledge	50	2.04	1.52	49	-9.08	<.001	-1.28
Awareness	50	1.78	1.29	49	-12.09	<.001	-1.72
<i>Hockey</i>							

² Here we report partial Eta squared (η^2) which is a measure of the amount of variance in the dependent variable explained by a given independent variable plus its associated error variance. Following Ellis (2010), values greater than or equal to .14 are large effects, greater than or equal to .06 but less than .14 are medium effects, and greater than or equal to .01 but less than .06 are small effects (see also Rose, 2015; Rose, 2017; Rose and Schaffer, 2017; Rose, Tobia and Schaffer, 2019; Rose, Machery, Stich, et al., 2020; Turri, Rose and Buckwalter, 2018; Nadelhoffer, Rose, Buckwalter and Nichols, forthcoming; Rose, Buckwalter and Nichols, 2017; Turri, Buckwalter and Rose, 2016; Rose, Buckwalter and Turri, 2014, McLaughlin and Rose, 2018)

Intentionality	50	5.40	1.79	49	5.51	<.001	0.78
Knowledge	50	2.10	1.50	49	-8.94	<.001	-1.26
Awareness	50	1.84	1.39	49	-10.98	<.001	-1.55
<i>Pool</i>							
Intentionality	50	5.28	1.84	49	4.91	<.001	0.69
Knowledge	50	2.68	1.71	49	-5.42	<.001	-0.77
Awareness	50	2.26	1.72	49	-7.13	<.001	-0.77
<i>Air Hockey</i>							
Intentionality	50	5.02	1.82	49	3.95	<.001	0.56
Knowledge	50	2.42	1.86	49	-5.99	<.001	-0.084
Awareness	50	2.46	2.04	49	-5.33	<.001	-0.75

Results for each measure in each scenario are shown in Figure 1.



Figure 1: Mean ratings for judgment type across scenarios with 95% confidence intervals

In cases involving skilled action, participants were overwhelmingly inclined to say that agents intentionally performed the actions we asked about—e.g., sinking a shot—despite the fact that the agents failed to satisfy the knowledge and awareness conditions regarding those actions. Are these

findings restricted to only highly skilled actions? Might habitual actions, which require very little skill, be viewed differently?

2.3. Study 2

To address whether habitual actions are viewed differently than skilled actions, we conducted another study. We again recruited two hundred and fifty people using the same procedure as in experiment 1. Participants were randomly assigned to one of five scenarios. As in experiment 1, participants first read a brief scenario and then responded to three test statements.

Here is one scenario (see Appendix for full set of scenarios):

[Door] Suzy locks her door every morning as she leaves for work. On her way out to work one morning, she locks the door. But, because she is preoccupied with thoughts about her day, she is completely unaware of doing so. She gets to her car, pauses, and wonders if she locked the door. Because she was unaware of locking it, she didn't even remember locking it. So she walks back from her car to check if she locked the door.

Beneath the scenario, on the same page, participants responded to three test statements (presented in random order):

When Suzy was locking her door, she knew that she was locking it. (Knowledge)

While Suzy was locking her door, she was aware of locking it. (Awareness)

Suzy intentionally locked her door. (Intentionality)

Responses to all test statements were collected on the same 7-point Likert scale used in experiment 1.

2.4. Results for Study 2

We conducted a Mixed ANOVA with Question (Knowledge, Awareness, Intentionality) as a within subject factor and Scenario (Door, Coffee, Heater, Oven, Kiss) as between subject factor. There was a main effect of Question, $F(2, 490)=133.343$, $p<.001$, $\eta^2=.352$ but no effect of Scenario, $F(1,$

245)=.575, $p=.681$, $\eta^2=.009$ and no Question by Scenario interaction, $F(8, 490)=1.298$, $p=.242$, $\eta^2=.021$. Importantly, pairwise comparisons—and again using a Bonferonni correction of .016—indicated that Knowledge ($M=3.32$, $SD=1.82$) was significantly different from Awareness ($M=3.05$, $SD=1.81$), $p=.009$ and Intentionality ($M=4.88$, $SD=1.74$), $p<.001$, and that Awareness was significantly different from Intentionality, $p<.001$.

One sample t-tests revealed that mean intentionality attribution was significantly above the midpoint in all scenarios while mean knowledge attributions, except in the oven scenario, were significantly below the midpoint in all scenarios. Awareness attributions were significantly below the midpoint in all scenarios (see Table 2).

Table 2

Experiment 2. One sample t-tests. Test value = 4

Measure	N	M	SD	df	t	p	d
<i>Door</i>							
Intentionality	50	5.20	1.61	49	5.25	<.001	0.74
Knowledge	50	3.10	1.65	49	-3.84	<.001	-0.54
Awareness	50	2.92	1.77	49	-4.31	<.001	-0.61
<i>Coffee</i>							
Intentionality	51	5.17	1.45	50	5.78	<.001	0.81
Knowledge	51	3.43	1.86	50	-2.17	.034	-0.31
Awareness	51	3.29	1.91	50	-2.63	.011	-0.37
<i>Heater</i>							
Intentionality	50	4.78	1.63	49	3.37	<.001	0.49
Knowledge	50	3.30	1.70	49	-2.90	.006	-.041
Awareness	50	3.24	1.82	49	-2.94	.005	-0.42
<i>Oven</i>							
Intentionality	50	4.58	1.95	49	2.09	.042	0.29
Knowledge	50	3.52	1.97	49	-1.72	.091	-0.24
Awareness	50	3.06	1.95	49	-3.40	<.001	-0.48
<i>Kiss</i>							
Intentionality	49	4.67	1.96	48	2.40	.020	0.34
Knowledge	49	3.22	1.91	48	-2.83	.007	-.041
Awareness	49	2.73	1.55	48	-5.70	<.001	-0.82

Results for each measure in each scenario are shown in Figure 2.




Figure 2: Mean ratings for judgment type across scenarios with 95% confidence intervals

2.5. Discussion

Even when moving from cases involving skilled action to cases involving habitual action, we continue to find that a majority of participants say that an agent intentionally performed an action—e.g., locking a door—despite the fact that the agent failed to satisfy the knowledge and awareness conditions regarding that action. Our findings thus indicate that, on the ordinary view, neither knowledge nor awareness of doing something is required for doing it intentionally.

We find it surprising that quite a few participants gave a response of 4 or higher to the awareness or knowledge questions in our second experiment. Perhaps these respondents did not buy our stories. So suppose our questions about knowledge and awareness are treated as comprehension checks and we exclude participants who ascribed knowledge, awareness, or both. Removing those who gave a rating of 4 or higher on these conditions, the results are as follows: intentionality attributions for

those who denied awareness ($M=4.62$, $SD=1.88$) were significantly above the midpoint, $t(165)=4.20$, $p<.001$; intentionality attributions for those who denied knowledge ($M=4.43$, $SD=1.94$) were significantly above the midpoint, $t(147)=2.71$, $p=.007$; and intentionality attributions for those who denied both knowledge and awareness ($M=4.38$, $SD=1.96$) were also significantly above the midpoint, $t(132)=2.25$, $p=.026$. This is evidence that those who clearly deny knowledge, awareness or both nonetheless tend to say that an agent intentionally performed the action at issue.

3. Conclusion

Our aim was to evaluate KAT empirically. We found that majority responses to our vignettes are at odds with KAT. Our results show that, on an ordinary view of matters, neither knowledge nor awareness of doing something is necessary for doing it intentionally. We tested cases of skilled action and habitual action, and we found that, for both, people ascribed intentionality to an action at an appreciably higher rate than knowledge and awareness.

Some proponents of KAT might maintain that our results have no bearing on KAT. They may claim that folk intuitions are often mistaken because these intuitions are not the results of the kind of rigorous thinking philosophers utilize to elucidate concepts (see e.g., Deutsch 2009; Hales 2006; Ludwig 2007; Williamson 2005, 2011). Perhaps. But, at least in this case, it is not at all clear what might be leading people astray. It is thus doubtful that they are succumbing to some sort of error. If folk intuitions in this case aren't plausibly explained as the result of some error, then proponents of KAT should find our results worrisome. If fellow competent speakers of English have intuitions that conflict with those of some philosophers, this raises a red flag. This is because the concepts that philosophers care to elucidate stem from folk concepts. Moreover, many philosophers are concerned with what actually constitutes the concept of intentional action that people naturally use in everyday parlance. If theorists stray too far from what the folk naturally take or do not take to be an intentional action, they may be analyzing a different concept from the one they were interested in analyzing at the

outset. If so, they may be viewed as simply changing the subject (see Mele 2001: 27). Of course, not all philosophers are concerned with providing an analysis of intentional action that respects the folk view. But insofar as the folk view is taken to be relevant, our findings put pressure on KAT.

We grant that KAT cannot be rejected out of hand simply on the basis of our results. The fact that our respondents seem to disagree with an underlying conception of intentional action when considering certain cases does not strike a toppling blow to KAT. But the fact remains that our results have lifted a corner of the theoretical fabric to reveal an inconsistency between the intuitions of non-philosophers and KAT. We conclude that the presence of this inconsistency is significant and cannot go unanswered by KAT's proponents. It remains for them to explain how these results can peacefully coexist with KAT's truth.

References

- Alicke, M., & Rose, D. 2010. Culpable control or moral concepts? *Behavioral and Brain Sciences*, 33(4): 330-331.
- Alicke, M. D., & Rose, D. 2012. Culpable control and causal deviance. *Social and Personality Psychology Compass*, 6(10): 723-735.
- Alicke, M. D., Rose, D., & Bloom, D. 2011. Causation, norm violation, and culpable control. *The Journal of Philosophy*, 108(12): 670-696.
- Anscombe, G. E. M. 1957. *Intention*. Ithaca, NY: Cornell University Press.
- Baayen, R. H., Davidson, D. J., and Bates, D. M. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language* 59 (4): 390-412.
- Clark, H. H. 1973. The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior* 12 (4): 335-59.
- Davidson, D. 1963. Actions, reasons and causes. *Journal of Philosophy* 60(23): 685-700.
- Deutsch, M. 2009. Experimental philosophy and the theory of reference. *Mind & Language* 24(4): 445-466.
- Davidson, D. 1980. *Agency. Essays on Actions and Events*. Oxford University Press.
- Ellis, P. 2010. *The Essential Guide to Effect Sizes: Statistical Power, Meta-Analysis and the Interpretation of Research Results*. Cambridge University Press.
- Goldman, A. 1970. *A Theory of Human Action* Princeton, New Jersey: Princeton University Press.
- Gorr, M. & Horgan, T. 1982. Intentional and unintentional actions. *Philosophical Studies* 41(2): 251-262.
- Hales, S. D. 2006. *Relativism and the Foundations of Philosophy* Cambridge, Mass.: MIT Press.
- Hampshire, S. 1982. *Thought and Action* London: Chatto and Windus.

- Judd, C. M., Westfall, J., & Kenny, D. A. 2012. Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology* 103(1): 54-69.
- Knobe, J. 2003. Intentional action and side effects in ordinary language. *Analysis*, 63(3): 190-194.
- Lanteri, A. 2009. Judgements of intentionality and moral worth: Experimental challenges to Hindriks. *The Philosophical Quarterly* 59(237): 713–720.
- Ludwig, K. 2007. The epistemology of thought experiments: First person versus third person approaches. *Midwest Studies in Philosophy* 31(1): 128-159.
- Marcus, E. 2019. Reconciling practical knowledge with self-deception. *Mind* 128(512): 1205–1225.
- McLaughlin, B., & Rose, D. 2018. On the matter of robot minds. In T. Lombrozo, S. Nichols & J. Knobe (Eds.), *Oxford Studies in Experimental Philosophy, Volume 2* (pp. 316-332). Oxford University Press.
- Mele, A. 2001. Acting intentionally: probing folk notions. In B. Malle, L. Moses, and D. Baldwin, eds. *Intentions and Intentionality: Foundations of Social Cognition*, 27-43 Cambridge, Mass.: MIT Press.
- Nadelhoffer, T. 2004. The Butler problem revisited. *Analysis* 64 (3): 277-284.
- Nadelhoffer, T., Rose, D., Buckwalter, W., & Nichols, S. forthcoming. Natural Compatibilism, indeterminism and intrusive metaphysics. *Cognitive Science*.
- Olsen, C. 1969. Knowledge of one's own intentional actions. *Philosophical Quarterly* 19(77): 324-336.
- Rose, D. 2015. Persistence through function preservation. *Synthese*, 192(1): 97-146.
- Rose, D. 2017. Folk intuitions of actual causation: a two-pronged debunking explanation. *Philosophical Studies*, 174(5): 1323-1361.
- Rose, D., & Schaffer, J. 2017. Folk mereology is teleological. *Noûs*, 51(2): 238-270.

- Rose, D., Buckwalter, W., & Nichols, S. 2017. Neuroscientific prediction and the intrusion of intuitive metaphysics. *Cognitive Science*, 41(2): 482-502.
- Rose, D., Buckwalter, W., & Turri, J. 2014. When words speak louder than actions: Delusion, belief, and the power of assertion. *Australasian Journal of Philosophy*, 92(4): 683-700.
- Rose, D., Schaffer, J., & Tobia, K. (forthcoming). Folk teleology drives persistence judgments. *Synthese*, 1-19.
- Rose, D., Machery, E., Stich, S., Alai, M., Adriano, A., Berniunas, R., Buchtel, E., Chatterjee, A., Cheon, H., Cho, I., Cohnitz, D., Cova, F., Dranseika, V., Eraña Lagos, A., Ghadakpour, L., Grinberg, M., Hannikainen, I., Hashimoto, T., Horowitz, A., Hristova, E., Jraissati, Y., Kadreva, V., Karasawa, K., Kim, H., Kim, Y., Lee, M., Mauro, C., Mizumoto, M., Moruzzi, S., Olivola, C., Ornelas, J., Osimani, B., Rosas Lopez, A., Romero, C., Sangoi, M., Sereni, A., Songhorian, S., Sousa, P., Struchiner, N., Tripodi, V., Usui, N., Vázquez del Mercado, A., Volpe, G., Vosgerichian, H., Zhang, X., & Zhu, J. forthcoming. The ship of Theseus puzzle. In T. Lombrozo, S. Nichols & J. Knobe (eds.), *Oxford Studies in Experimental Philosophy*, Volume 3. Oxford University Press.
- Schwenkler, J. 2019. *Anscombe's Intention: A Guide*. Oxford University Press.
- Turri, J., Buckwalter, W., & Rose, D. 2016. Actionability judgments cause knowledge judgments. *Thought: A Journal of Philosophy*, 5(3): 212-222.
- Turri, J., Rose, D., & Buckwalter, W. 2018. Choosing and refusing: doxastic voluntarism and folk psychology. *Philosophical Studies*, 175(10): 2507-2537.
- Williamson, T. 2005. Armchair philosophy, metaphysical modality and counterfactual thinking. *Proceedings of the Aristotelian Society* 105: 1-23.
- Williamson, T. 2011. Philosophical expertise and the burden of proof. *Metaphilosophy* 42(3): 215-229.

Appendix

Experiment 1

[Golf] Andy makes a hole-in-one 92% of the time when playing putt-putt. He is at the indoor Putt-Putt Fun Center one evening, practicing his shot as he usually does. He lines up his first shot at the first hole. Just as he swings the putter and hits the ball, his friend Suzy plays a trick on him: she throws a black bag over his head and screams at the top of her lungs. Due to this, Andy could not see or even hear whether he made the shot. He is completely unaware of whether he sank the shot. But he did in fact sink the shot.

Please indicate the extent to which you agree with the following statements:

When Andy was sinking the shot, he knew that he was sinking it.

While Andy was sinking the shot, he was aware of sinking the shot.

Andy intentionally sank the shot.

[Hockey] Andy has been practicing for an upcoming hockey contest. To win the contest, he needs to shoot a puck 50 feet into a bucket. This is called “scoring a goal” in this contest. He is able to do this 92% of the time. On the day of the contest, he lines up to take the shot. Just as he hits the puck, a drunken fan accidentally hits him in the head with a beer bottle. It momentarily knocks Andy unconscious. Due to this, Andy could not see or even hear whether he scored a goal. He is completely unaware of whether he scored a goal. But he did in fact score one.

Please indicate the extent to which you agree with the following statements:

When Andy was scoring the goal, he knew that he was scoring it.

While Andy was scoring the goal, he was aware of scoring the goal.

Andy intentionally scored the goal.

[Pool] Andy sinks 92% of his attempts at a certain trick shot when playing pool. To impress some people at the pool hall, he lines up his shot. Just as he hits the cue ball, one of the women he is trying to impress jumps in front of him and blows an air horn in his face. Due to this, Andy could not see or even hear whether he made the shot. He is completely unaware of whether he sank the shot. But he did in fact sink the shot.

Please indicate the extent to which you agree with the following statements:

When Andy was sinking the shot, he knew that he was sinking it.

While Andy was sinking the shot, he was aware of sinking the shot.

Andy intentionally sank the shot.

[Air Hockey] Andy scores on 92% of his shots when practicing air hockey. He is at the arcade one evening practicing some air hockey. He lines up his shot. Just as he hits the air hockey puck, a kid who had been shaking up a soda opens it and soda sprays all over Andy. It gets all over his clothes and even in his eyes. The kid's mom screams at the top of her lungs. Due to this, Andy could not see or even hear whether he made the shot. He is completely unaware of whether he scored a goal. But he did in fact score a goal.

Please indicate the extent to which you agree with the following statements:

When Andy was scoring the goal, he knew that he was scoring it.

While Andy was scoring the goal, he was aware of scoring the goal.

Andy intentionally scored the goal.

Experiment 2

[Coffee] Suzy turns her coffee machine off every morning before she takes a shower. One morning before she gets in the shower, she turns her coffee machine off. But, because she is preoccupied with

thoughts about her day, she is completely unaware of doing so. She gets in the shower and half-way through wonders if she turned the coffee machine off. Because she was unaware of turning it off, she doesn't even remember turning it off. So she gets out of the shower to check if she turned the coffee machine off.

Please indicate the extent to which you agree with the following statements:

When Suzy was turning the coffee machine off, she knew that she was turning it off.

While Suzy was turning the coffee machine off, she was aware of turning it off.

Suzy intentionally turned off her coffee machine.

[Heater] Suzy turns her heater off every day before she leaves her office. One evening, she turns her heater off before leaving her office. But, because she is preoccupied with thoughts about her day, she is completely unaware of doing so. As she leaves the building she begins to wonder if she turned her heater off. Because she was unaware of turning it off, she doesn't even remember turning it off. So she walks back to her office to check if she turned her heater off.

Please indicate the extent to which you agree with the following statements:

When Suzy was turning the heater off, she knew that she was turning it off.

While Suzy was turning the heater off, she was aware of turning it off.

Suzy intentionally turned off her heater.

[Oven] Suzy loves to bake. Every day when she gets home from work, the first thing she does is turn her oven on. One evening, she returns home from work and turns her oven on. But, because she is preoccupied with thoughts about her day, she is completely unaware of doing so. She then goes to the living room to watch the news. As she is watching tv, she starts to wonder if she turned her oven on.

Because she was unaware of turning it on, she doesn't even remember turning it on. So she walks back to the kitchen to check if she turned the oven on.

Please indicate the extent to which you agree with the following statements:

When Suzy was turning the oven on, she knew that she was turning it on.

While Suzy was turning the oven on, she was aware of turning it on.

Suzy intentionally turned the oven on.

[Kiss] Suzy drops her son off at preschool every morning on the way to work. When she walks him in, she always gives him a kiss on the cheek before saying goodbye. One morning, she drops her son off at preschool and kisses him on the cheek before leaving. But, because she is preoccupied with thoughts about her day, she is completely unaware of doing so. As she leaving the building, she starts to wonder if she kissed her son on the cheek. Because she was unaware of kissing him on the cheek, she doesn't even remember kissing him on the cheek. So she walks back into the preschool to ask him whether she kissed him on the cheek.

Please indicate the extent to which you agree with the following statements:

When Suzy was kissing her son on the cheek, she knew that she was kissing him on the cheek. While

Suzy was kissing her son on the cheek, she was aware of kissing him on the cheek.

Suzy intentionally kissed her son on the cheek.