

[A] Salience and Epistemic Egocentrism: An Empirical Study

Joshua Alexander, Chad Gonnerman, and John Waterman¹

[A] Introduction

Jennifer Nagel (2010) has recently proposed a fascinating account of the decreased tendency to attribute knowledge in conversational contexts in which unrealized possibilities of error have been mentioned. Her account appeals to epistemic egocentrism, or what is sometimes called the curse of knowledge, an egocentric bias to attribute our own mental states to other people (and sometimes our own future and past selves). When asked to consider a story in which someone is making a judgment, and told that it's possible that her judgment is wrong, we treat our concerns as hers and penalize her for failing to appropriately respond to those concerns.

This account of the relationship between our willingness to attribute knowledge and what possibilities have been made salient in a given conversational context may be particularly appealing to invariantists. It provides them with a way to explain why, although standards for knowledge attribution don't depend on what possibilities have been made salient in a given conversational context, we are nevertheless less willing to attribute knowledge in conversational contexts that include mention of the possibility that the subject might be wrong than in contexts in which no such possibilities are raised.

Our aim in this paper is to investigate the empirical merits of Nagel's hypothesis about the psychology involved in knowledge attribution. We begin by briefly reviewing the epistemological and psychological background. After setting the stage, we present four new studies showing that our willingness to attribute knowledge is sensitive to what possibilities have been made salient in a given conversational context, that this sensitivity can be, at least in part,

explained in terms of epistemic egocentrism, and that increased motivation doesn't seem to drive down our tendency to mistakenly project our own mental states onto others.² We conclude by previewing three additional studies involving individual differences, social distance, and a specific kind of interventional debiasing strategy.

[A] Section One: Philosophical and Psychological Background

One of the most hotly debated questions in contemporary epistemology is whether standards for knowledge attribution depend on what possibilities have been made salient in a given conversational context.³ Contextualists contend that it does, arguing that whether it is right to say that someone knows something depends at least in part on whether our conversational context includes any mention of the possibility that she might be wrong (DeRose, 1992; Cohen, 1999). When a conversational context includes mention of the possibility that the subject might be wrong, the standards for knowledge attribution rise; the verb “to know” comes to refer to a more stringent relation, and subjects who could once have been truly described as knowing something might no longer be truly described that way. Invariantists contend that it does not, arguing that what possibilities are mentioned in a given conversational context does not affect whether someone knows something; in their view, the verb “to know” always denotes the same relation (Hawthorne, 2004; Williamson, 2005; Nagel, 2010).

What makes this debate especially interesting is that both sides agree that our willingness to attribute knowledge seems to depend at least in part on what possibilities have been made salient in a given conversational context: we seem to be more willing to say that someone knows something in conversational contexts that don't include any mention of the possibility that she

might be wrong than in contexts that do. Consider, for example, the following two vignettes (Nagel 2010):

Plain Story

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. Q: Does he know that the table is red?

More Detailed Story

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. Q: Does he know that the table is red?

Contextualists and invariantists agree that we are more willing to say that John Doe knows that the table is red in the plain story than we are to say that John Doe knows that the table is red in the more detailed story because they agree that we become less willing to attribute knowledge in conversational contexts that include any mention of possible error. This creates a challenge for invariantists, however. If invariantists are right, and standards for knowledge attribution don't depend on what possibilities have been made salient in a given conversational context, then they need some way of explaining why we are nevertheless more willing to attribute knowledge in conversational contexts that don't include any mention of the possibility that the subject might be wrong than in contexts that do.

Jennifer Nagel (2010) has recently offered a fascinating psychological explanation of the relationship between our willingness to attribute knowledge and what possibilities have been made salient in a given conversational context.⁴ Nagel's account is grounded in the idea that we have a difficult time representing perspectives more naïve than our own and, in particular, we struggle to suppress privileged information when evaluating other people's judgments (Nickerson, 1999; Royzman et al., 2003; Birch & Bloom, 2004, 2007). This psychological bias, known as epistemic egocentrism or the curse of knowledge, explains why we seem to be less willing to attribute knowledge in conversational contexts that include any mention of the possibility of error: we evaluate other people's judgments as though they share our privileged information and subsequently penalize them for failing to respond to that information in the way we think they should.⁵ In other words, and in terms of the two vignettes just described, we are less willing to say that John Doe knows that the table is red in the more detailed story than we are in the plain story precisely because we treat him in the detailed story as sharing our concerns and failing to act accordingly.

Nagel's argument depends on specific empirical claims about the psychology involved knowledge attribution, claims that can be tested, and our plan in the next section is to examine the empirical merits of Nagel's account.

[A] Section Two: Our Studies

[B] Study 1: Salience Effects

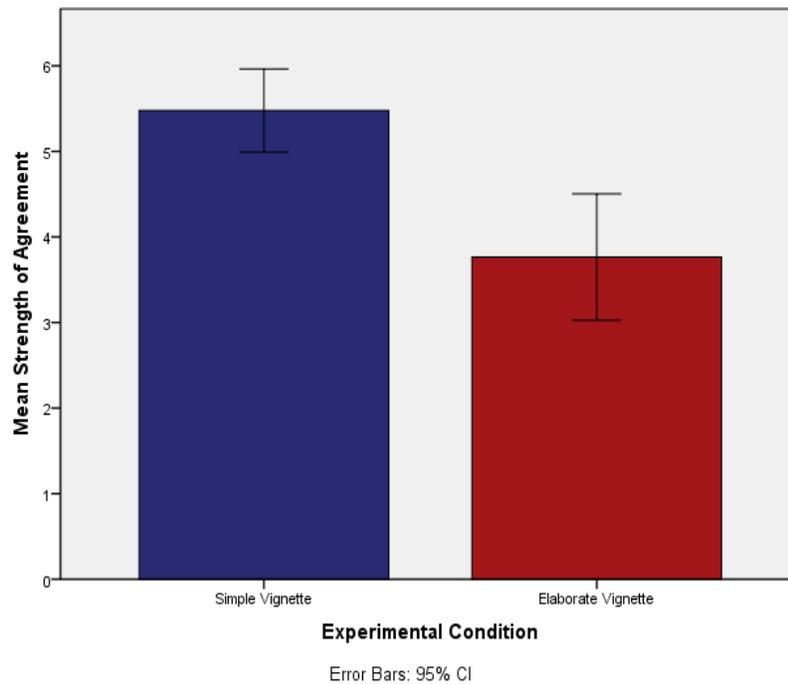
Contextualists and invariantists agree that we are less willing to attribute knowledge in conversational contexts that include mention of unrealized possibilities of error, but disagree about how to explain this. As we might put it, contextualists and invariantists agree that salience

matters, but disagree about why. Somewhat surprisingly, the results of several early empirical studies suggested that salience doesn't matter. These studies seemed to show that people are just as willing to say that someone knows something when the possibility of being wrong has been made salient as they are to say that someone knows something when that possibility goes unmentioned. The studies focused on several different versions of Keith DeRose's (1992, 2009) famous bank cases, which involve a husband and wife stopping by a bank on their way home from work on a Friday afternoon intending to deposit their paychecks. When they notice that the lines inside are quite long, the husband suggests coming back the following day to make the deposits, noting to his wife that he knows that the bank will be open on Saturday because he was there two Saturdays ago. The relevant difference between the two cases, at least for our purposes, is that, in the second case, the wife mentions the possibility that her husband might be wrong, noting that banks sometimes change their hours. Contextualists and invariantists agree that folk knowledge attributions will track this difference, that people will be less inclined to say that the husband knows that the bank will be open on Saturday when the possibility of being wrong has been made salient than when that possibility goes unmentioned. In contrast, what the early studies found was that people are just as willing to say that the husband knows that the bank will be open on Saturday when the possibility that he might be wrong has been made salient as they are to say that he knows that the bank will be open on Saturday when that possibility goes unmentioned (Buckwalter, 2010; Feltz & Zerpentine, 2010; May et al., 2010).

While these studies give us reason to worry that salience might not matter, several recent objections warn against drawing this conclusion too quickly. One worry with the early studies is that they failed to make sufficiently salient the possibility that the bank will be closed on Saturday: merely mentioning a possibility doesn't not necessarily make that possibility salient,

particularly when that possibility seems strange or improbable (Schaffer & Knobe, 2011). And, in fact, it does seem that people are less inclined to say that the husband knows that the bank will be open on Saturday when more care is given to making that possibility salient – for instance by embedding it in the context of a personal anecdote (Schaffer & Knobe, 2011; Buckwalter, this volume). A second worry with the early studies is that several of them involved asking participants to make knowledge attributions rather than to evaluate knowledge attributions, and that when participants were asked to evaluate knowledge attributions they were asked to do so in situations where it was natural to deny knowledge (DeRose, 2011). And, when adjustments are made to address these worries, it again seems that people are less inclined to say that the husband knows that the bank will be open on Saturday when the possibility of being wrong has been made salient than when that possibility goes unmentioned (Buckwalter, this volume).⁶

In light of this unsettled experimental landscape, another look at salience outside the realm of bank cases seemed to us to be in order. To do this, we constructed a study using Nagel’s plain story and more detailed story. Participants ($N = 40$, Age $M = 30$, Female = 28%) were recruited via Amazon Mechanical Turk.⁷ Participants received one of the vignettes, and were then asked to indicate the extent to which they agreed or disagreed with the claim that John knows that the table is red. Answers were assessed using a six-point Likert scale with 1 = strongly disagree and 6 = strongly agree. Contrary to the results of the early studies, but consistent with the results of more recent studies of the influence of salience on folk knowledge attribution, we found that participants were significantly more willing to attribute knowledge in the simple story ($M = 5.50$, $SD = 1.14$) than in the more detailed story ($M = 3.78$, $SD = 1.40$). A planned comparison showed statistically significant difference ($t = 4.29$, $df = 38$, $p < .001$). Cohen’s $d = 1.36$, a “large” effect. These results can be visualized as follows:



By focusing on vignettes where the possibility of error is made sufficiently salient, these results contribute to the recent trend of studies that have found a salience effect. But while these results suggest that salience does influence folk knowledge attributions, neither this study nor the ones that preceded it suggest why salience influences folk knowledge attributions. Our next step was to focus on this question.

[B] Study 2: Epistemic Egocentrism and the Curse of Knowledge, Round 1

Epistemic egocentrism involves misrepresenting other people's mental states. For instance, we might mistakenly treat our concerns as theirs, and penalize them for failing to appropriately respond to these concerns. This may help to explain why we are less willing to say

that John Doe knows that the table is red in the more detailed story than we are in the plain story. In the more detailed story, we treat him as sharing our concerns about the possibility of abnormal lighting conditions and subsequently penalize him for failing to check whether the lighting is normal.

Epistemic egocentrism involves two important empirical predictions. First, it predicts that there should not be a significant difference between how we assess cases where the relevant information is privileged and how we assess cases where it is not. That is, epistemic egocentrism predicts that there should be relatively little difference between our assessment of narrator cases, cases where the relevant information is shared only with the reader, and our assessment of subject cases, cases where the relevant information is shared with both reader and subject.⁸ Second, epistemic egocentrism predicts that there should be relatively little difference between how we assess entertain cases, cases where the subject is portrayed as entertaining the possibility that she is wrong, and how we assess neutral cases, cases that leave open whether or not the subject is entertaining the possibility that she is wrong.

In order to test these two predictions, we constructed the following four vignettes, staying as close as possible to Nagel's more detailed case:

Narrator Defeater (Neutral)

John and Mary are in a furniture store. John is looking at a bright red table under normal lighting conditions. He believes the table is red. However, a white table under red lighting would look exactly the same. John has not checked whether the lighting is normal, or whether there might be a red spotlight shining on the table.

Narrator Defeater (Entertain)

John and Mary are in a furniture store. John is looking at a bright red table under normal lighting conditions. He believes the table is red. However, a white table under red lighting would look exactly the same. John thinks about this, but does not check whether the lighting is normal, or whether there might be a red spotlight shining on the table.

Subject Defeater (Neutral)

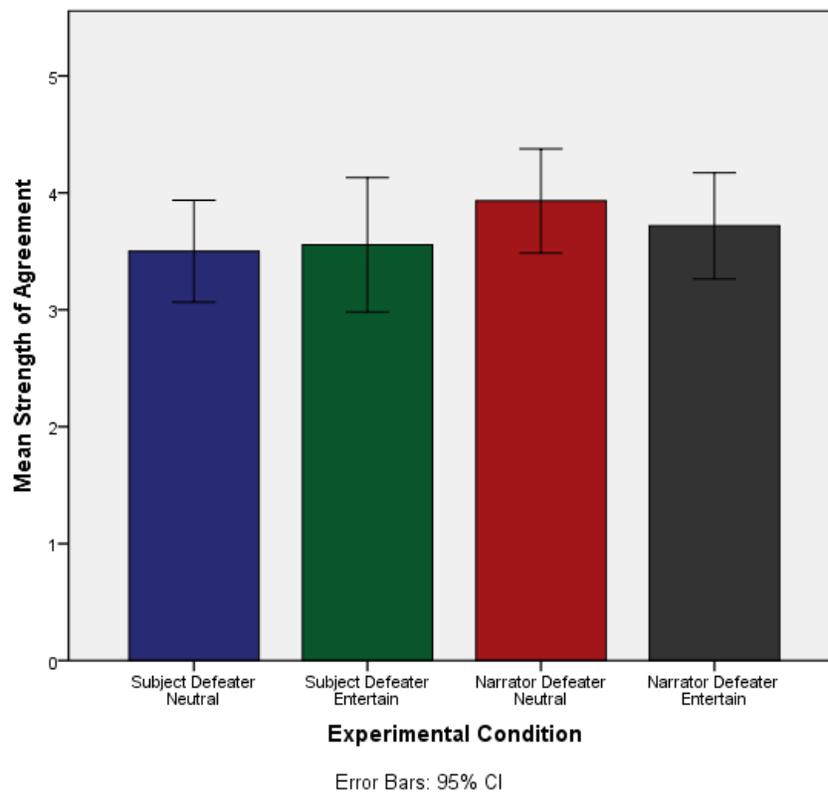
John and Mary are in a furniture store. John is looking at a bright red table under normal lighting conditions. He believes the table is red. Mary points out, however, that a white table under red lighting would look exactly the same. John has not check whether the lighting is normal, or whether there might be a red spotlight shining on the table.

Subject Defeater (Entertain)

John and Mary are in furniture store. John is looking at a bright red table under normal lighting conditions. He believes the table is red. Mary points out, however, that a white table under red lighting would look exactly the same. John thinks about this, but does not check whether the lighting is normal, or whether there might be a red spotlight shining on the table.

Participants ($N = 187$, Age $M = 27$, Females 26%) received one of these vignettes, and were then asked to indicate the extent to which they agreed or disagreed with the claim that John knows

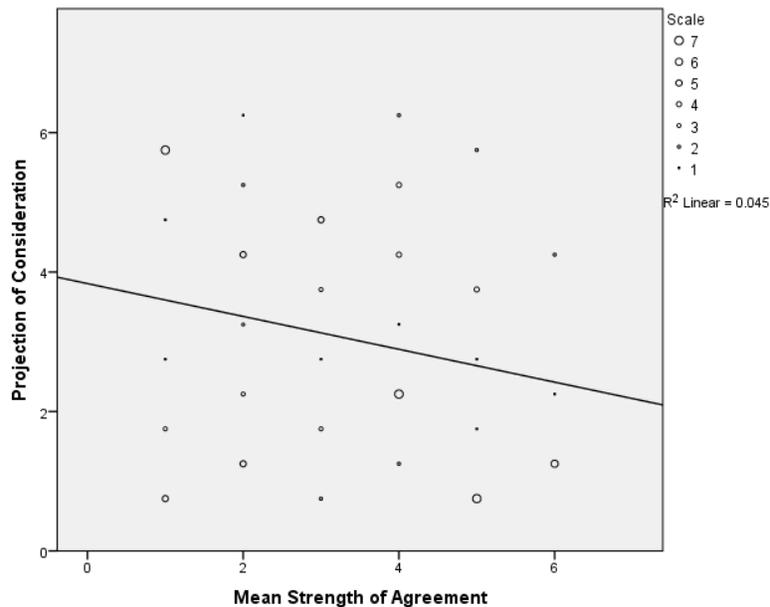
that the table is red. Answers were assessed using a six-point Likert scale with 1 = strongly disagree and 6 = strongly agree. As expected, we found that there was no statistically significant difference between how participants assessed narrator cases and how they assessed subject cases, nor was there a statistically significant difference between how participants assessed entertain cases and how they assessed neutral cases. The means and standard deviations for the four conditions were: Narrator Defeater Neutral ($n = 58$, $M = 3.93$, $SD = 1.70$), Narrator Defeater Entertain ($n = 45$, $M = 3.78$, $SD = 1.49$), Subject Defeater Neutral ($n = 48$, $M = 3.5$, $SD = 1.50$), Subject Defeater Entertain ($n = 36$, $M = 3.56$, $SD = 1.70$). The results of a 4x1 ANOVA showed no interaction between the various conditions ($F = 0.783$, $p = 0.505$).⁹ A subsequent pair-wise comparison failed to find a statistically significant difference between entertain and neutral conditions in either subject cases ($t = -.159$, $df = 82$, $p = .874$) or in narrator cases ($t = .479$, $df = 101$, $p = .633$). These results can be visualized as follows:



These results give us some reason to think that something like epistemic egocentrism is driving our evaluation of other people's judgments, at least in some conversational contexts that include mention of unrealized possibilities of error. Since we are treating our concerns as theirs, it doesn't matter whether they are portrayed as sharing these concerns. It doesn't even matter whether they are portrayed as being aware of them at all. All that seems to matter is that they haven't done enough to appropriately respond to these concerns. And, this is just what we'd expect if epistemic egocentrism were playing a role in our evaluations of other people's judgments.

[B] Study 3: Epistemic Egocentrism and the Curse of Knowledge, Round 2

It turns out that there is another way to test this. If we are projecting our concerns onto others and penalizing them for failing to appropriately respond to these concerns, then we should find a negative correlation between the projection of our concerns onto others and our willingness to attribute knowledge to them, at least in situations where they fail to appropriately respond to these concerns. In order to test this prediction, we gave participants ($N = 93$, Age $M = 28$, Female = 32%) Nagel's more detailed case, and then asked them to indicate the extent to which they agreed or disagreed with the claim that John knows that the table is red and the claim that John is considering the possibility that he's looking at a white table under a red spotlight. As before, answers were assessed using a six-point Likert scale with 1 = strongly disagree and 6 = strongly agree. We found the expected negative correlation ($r = -.211$, $p = .042$).¹⁰ The more likely participants were to say that John shared their concerns, the less likely they were to attribute knowledge to him, and vice versa. These results can be visualized as follows:



As the scatterplot helps to show, the inverse correlation that we found is a modest one. With r equaling -0.211 , we can only say that about 4% of the variability in responses to the knowledge probe is directly predictable from variability in responses to the considering probe, leaving the other 96% to other factors. At this point, it is difficult to determine precisely what is going on here. One possibility is that this modest correlation is precisely what we should expect, on Nagel's hypothesis.¹¹ After all, according to it, when we project our concerns onto subjects like John Doe, our projections are due to the operation of an unconscious cognitive bias. So we might well expect our participants to have rather limited access to not only the cause of their projection but perhaps to their projections as well, resulting in the modest correlation observed in our study. If this is right, then the key to improving our predictive powers, while working with Nagel's hypothesis, would be to find a better way of measuring the largely unconscious projections.

But there are other possibilities to consider. It could be that participants do in fact have reasonably decent access to their projections (though probably not the causes of their projections). If so, then to fully predict the variability we observed in the knowledge attributions of our participants, we would have to turn to factors other than those having to do with projective tendencies. The question then becomes, what are those other factors? Insofar as the dialectic between invariantists and contextualists is concerned, one rather strange possibility is that perhaps in conversational contexts that include mention of an unrealized possibility of error, participants do raise the standards for knowledge attribution, but not in a uniform way across all participants. That might be a welcomed result to the contextualists. For, when it comes to explaining salience effects, it seems that a partial victory for them would be victory enough. Of course, the results reported here don't settle the matter. But perhaps we can say this much: they do suggest that epistemic egocentrism is at least playing a role in our willingness to attribute knowledge in at least some cases that include mention of an unrealized possibility of error.

[B] Study 4: Salience, Epistemic Egocentrism, and Motivation

At this point, it might be natural to think that we just need to try harder to avoid mistakenly projecting our own mental states onto others. And some recent work in the social and cognitive sciences supports this kind of optimism, suggesting that certain kinds of epistemic egocentrism can be reduced with sufficient practice (Wu & Keysar, 2007; Cohen & Gunz, 2002), effort (Epley et al., 2004), or motivation (Epley et al., 2004). Yet, there is also reason to worry that our bias against adopting other perspectives cannot be cancelled completely. Hindsight bias and outcome bias, for example, turn out to be exceptionally resilient (Camerer et al., 1989; Krueger & Clement, 1994; Pohl & Hell, 1996), and motivation alone does not ensure that we

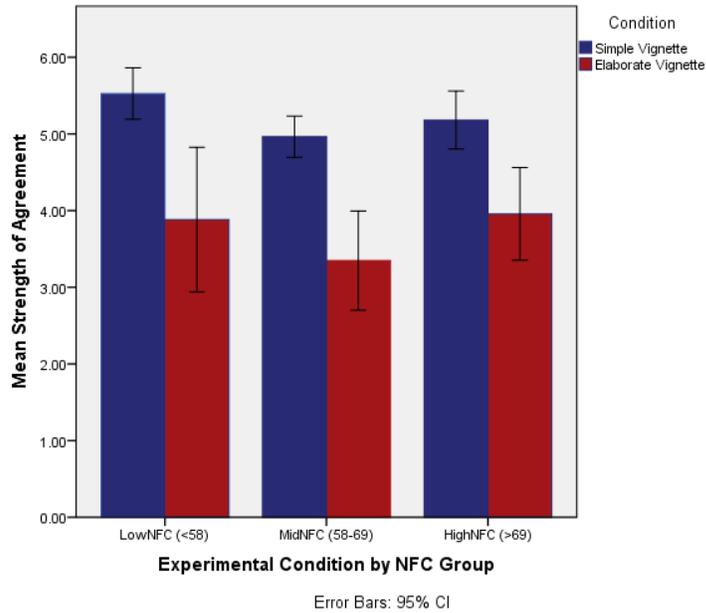
leave our own perspectives behind entirely (Epley & Caruso, 2009).

Against this background, we wanted to see whether sufficient motivation might help reduce epistemic egocentrism in the kinds of cases we've been discussing. Although motivation is often measured by means of financial incentive, we decided to measure it in terms of people's need for cognition (Cacioppo & Petty, 1982; for additional work on need for cognition and knowledge attributions, see Weinberg et al., 2012).¹² People's need for cognition corresponds to their intrinsic motivation to give a great deal of care and attention to cognitive tasks, and people with a high need for cognition have been shown to be less likely to be influenced by certain cognitive heuristics and biases (Priester & Petty, 1995; Cacioppo et al., 1996; Smith & Petty, 1996).¹³

Participants ($N = 126$, Age $M = 30$, Females 34%) received either Nagel's plain story or her more detailed story, were asked to indicate the extent to which they agreed or disagreed with the claim that John knows that the table is red (again, answers were assessed using a six-point Likert scale with 1 = strongly disagree and 6 = strongly agree), and were then asked to complete the need for cognition survey together with several additional demographic questions. The mean need for cognition (NFC) score for our participants was 63.3 ($SD = 13.4$). We defined NFC grouping by 1/2 standard deviation above and below that mean; NFC scores ranging from 18-57 were designated as "low NFC" ($n = 44$), those with NFC scores ranging from 58-69 were designated as "mid NFC" ($n = 48$), and those with NFC scores ranging from 70-90 were designated as "high NFC" ($n = 34$).

As before, participants seemed much more willing to attribute knowledge in the simple story than in the more detailed story. In the latter, participants seemed willing to penalize John Doe for failing to rule out the possibility that the lighting might be abnormal, a concern that was

shared only with them as readers of the vignette. The means and standard deviations for the two conditions were: Simple Story ($n = 68$; $M = 5.19$; $SD = 0.78$) and More Detailed Story ($n = 58$; $M = 3.69$, $SD = 1.54$). A planned comparison showed a statistically significant difference ($t = 7.073$, $df = 124$, $p < 0.001$, $d = 1.26$). What is particularly interesting in the current context is that increased motivation doesn't seem to drive down our tendency to mistakenly project our own mental states onto others in these kinds of cases: highly motivated people seem just as likely as highly unmotivated people to project their concerns onto John Doe.¹⁴ The means and standard deviations for the different conditions, breaking participants into three groups (high NFC, mid NFC, and low NFC), are as follows. High NFC: simple story ($n = 22$, $M = 5.18$, $SD = 0.85$), more detailed story ($n = 22$, $M = 3.91$, $SD = 1.41$). Planned comparison showed a statistically significant difference ($t = 3.621$, $df = 42$, $p = 0.001$, $d = 1.09$). Mid NFC: simple story ($n = 27$, $M = 4.96$, $SD = 0.71$), more detailed story ($n = 21$, $M = 3.29$, $SD = 1.52$). Planned comparison showed a statistically significant difference ($t = 5.079$, $df = 46$, $p < 0.001$, $d = 1.48$). Low NFC: simple story ($n = 19$, $M = 5.53$, $SD = 0.70$), more detailed story ($n = 15$, $M = 3.93$, $SD = 1.71$). Planned comparison showed a statistically significant difference ($t = 3.393$, $df = 17.674$, $p = 0.003$, $d = 1.28$). A two-way ANOVA showed that there was no significant interaction between NFC groupings and responses ($p = 0.70$).¹⁵ These results can be visualized as follows:



[A] Conclusion

We have seen that our willingness to attribute knowledge can be sensitive to what possibilities have been made salient in a given conversational context, that this sensitivity can, at least in part, be explained in terms of epistemic egocentrism, and that increased motivation doesn't seem to drive down our tendency to mistakenly project our own mental states onto others in the cases examined here. There is more to learn, and we think that three additional avenues of research will be particularly important going forward. Recent work in the social and cognitive sciences has shown that there are individual differences in our ability to set aside what we know (Kuhn, 1991; Stanovich & West, 1998; Musch & Wagner, 2007), and that social distance influences the tendency to project our own mental states onto other people (Robinson et al., 1995; Epley et al., 2004; and Ames, 2004). Studies have also shown that at least some kinds of epistemic egocentrism – hindsight bias, for example – can be overcome using a specific kind of interventional debiasing strategy known as the consider the opposite strategy (Lord et al., 1984;

Arkes et al., 1988; Arkes, 1991; Larrick, 2004). As we come to better understand the role that epistemic egocentrism plays in knowledge attribution, especially in conversational contexts involving privileged information, it will be important to study whether knowledge attributions track individual differences in susceptibility to the curse of knowledge, what role social distance plays in our ability to reason about what others know, and whether the influence of epistemic egocentrism on knowledge attribution can be mitigated by asking people to consider reasons why their initial judgments might have been wrong. These are projects for another day.

[A] Bibliography

Ames, D. (2004), 'Inside the mind reader's tool kit: projection and stereotyping in mental state inference'. Journal of Personality and Social Psychology, 87, 340-353.

Arkes, H. (1991), 'Costs and benefits of judgment errors: implications for debiasing'. Psychological Bulletin, 110, 486-498.

Arkes, H., Faust, D., Guilmette, T., and Hart, K. (1988), 'Eliminating the hindsight bias'. Journal of Applied Psychology, 73, 305-307.

Austin, J. L. (1990), Philosophical Papers (3rd Edition). Oxford: Oxford University Press.

Baron, J. and Hershey, J. (1988), 'Outcome bias in decision evaluation'. Journal of Personality and Social Psychology, 54, 569-579.

Birch, S. and Bloom, P. (2004), 'Understanding children's and adults' limitations in mental state reasoning'. Trends in Cognitive Sciences, 8, 255-260.

Birch, S. and Bloom, P. (2007), 'The curse of knowledge in reasoning about false beliefs'. Psychological Science, 18, 382-386.

Buckwalter, W. (2010), 'Knowledge isn't closed on Saturday'. Review of Philosophy and Psychology, 1, 395-406.

Buckwalter, W. (2012), 'Non-traditional factors in judgments about knowledge'. Philosophy Compass, 7, 278-289.

Buckwalter, W. (This volume), 'The mystery of stakes and error in ascriber intuitions', in J. Beebe (ed), Advances in Experimental Epistemology, London: Continuum Press.

Cacioppo, J. and Petty, R. (1982), 'The need for cognition'. Journal of Personality and Social Psychology, 42, 116-31.

Cacioppo, J., Petty, R., Feinstein, J. and Jarvis, B. (1996), 'Dispositional differences in cognitive motivation: the life and times of individuals varying in need for cognition'. Psychological Bulletin, 119, 197-253.

Cacioppo, J., Petty, R., and Kao, C. (1984), 'The efficient assessment of need for cognition'.

Journal of Personality Assessment, 48, 306-307.

Camerer, C., Loewenstein, G., and Weber, M. (1989), 'The curse of knowledge in economic settings'. Journal of Political Economy, 97, 1232-1254.

Cohen, S. (1999), 'Contextualism, skepticism, and the structure of reasons'. Philosophical Perspectives, 13, 587-89.

Cohen, D. and Gunz, A. (2002), 'As seen by the other... : Perspectives on the self in the memories and emotional perceptions of Easterners and Westerners'. Psychological Science, 13, 55-59.

DeRose, K. (1992), Contextualism and knowledge attributions. Philosophy and Phenomenological Research, 52, 913-929.

DeRose, K. (1999), 'Contextualism: an explanation and defense', in J. Greco and E. Sosa (eds) The Blackwell Guide to Epistemology. Oxford: Blackwell, pp. 187-205.

DeRose, K. (2009), The Case for Contextualism. Oxford: Oxford University Press.

DeRose, K. (2011), 'Contextualism, contrastivism, and X-Phi Surveys'. Philosophical Studies, 156, 81-110.

Epley, N., Keysar, B., Van Boven, L., and Gilovich, T. (2004), 'Perspective taking as egocentric anchoring and adjustment'. Journal of Personality and Social Psychology, 87, 327-339.

Epley, N. and Caruso, E. (2009), 'Perspective taking: Misstepping into others' shoes', in K. Markman, W. Klein, and J. Suhr (eds), Handbook of Imagination and Mental Simulation. New York: Psychology Press, pp. 295-309.

Fantl, J. and McGrath, M. (2010), Knowledge in an Uncertain World. Oxford: Oxford University Press.

Feltz, A. and Zarpentine, C. (2010), 'Do you know more when it matters less?' Philosophical Psychology, 23, 683-706.

Gilovich, T. (1991). How We Know What Isn't So. New York: The Free Press.

Goldman, A. I. (2006). Simulating Minds: The Philosophy, Psychology, and Neuroscience of Mindreading. Oxford: Oxford University Press.

Hawkins, S. and Hastie, R. (1990), 'Hindsight'. Psychological Bulletin, 107, 311-327.

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

Hawthorne, J. and Stanley, J. (2008), 'Knowledge and action'. Journal of Philosophy, 105, 571-590.

Holmes, D. S. (1968). 'Dimensions of projection'. Psychological Bulletin, 69, 248-268.

Krueger, J. and Clement, R. (1994), 'The truly false consensus effect'. Journal of Personality and Social Psychology, 67, 596-610.

Kuhn, D. (1991), The Skills of Argument. Cambridge: Cambridge University Press.

Larrick, R. (2004), 'Debiasing', in D. Koehler and N. Harvey (eds), Blackwell Handbook on Judgment and Decision Making. Oxford: Blackwell Publishing, pp. 316-338.

Lewis, D. (1996), 'Elusive knowledge'. Australasian Journal of Philosophy, 74, 549-567.

Lord, C., Lepper, M. and Preston, E. (1984), 'Considering the opposite: a corrective strategy for social judgment'. Journal of Personality and Social Psychology, 47, 1231-1243.

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

Musch, J., and Wagner, T. (2007), 'Did everybody know it all along? A review of individual differences in hindsight bias'. Social Cognition, 25, 64-82.

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

Nagel, J. (2010), 'Knowledge ascriptions and the psychological consequences of thinking about error'. The Philosophical Quarterly, 60, 286-306.

Nickerson, R. (1999), 'How we know – and sometimes misjudge – what others know'. Psychological Bulletin, 125, 737-759.

Pinillos, N. A. (2011), 'Some recent work in experimental epistemology'. Philosophy Compass, 6, 675-688.

Pohl, R. and Hell, W. (1996), 'No reduction of hindsight bias after complete information and repeated testing'. Organizational Behavior and Human Decision Processes, 67, 49-58.

Priester, J. and Petty, R. (1995), 'Source attributions and persuasion: perceived honesty as a determinant of message scrutiny'. Personality and Social Psychology Bulletin, 21, 637-54.

Robinson, R., Keltner, D., Ward, A., and Ross, L. (1995), 'Actual versus assumed differences in construal: "naïve realism" in intergroup perception and conflict'. Journal of Personality and

Social Psychology, 68, 404-417.

Royzman, E., Cassady, K., and Baron, J. (2003), 'I know, you know: epistemic egocentrism in children and adults'. Review of General Psychology, 7, 38-65.

Schaffer, J. and Knobe, J. (2010), 'Contrastive knowledge surveyed'. Nous, 46, 675-708.

Smith, S. and Petty, R. (1996), 'Message framing and persuasion: a message processing analysis'. Personality and Social Psychology Bulletin, 22, 257-68.

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

Stanovich, K. and West, R. (1998), 'Individual differences in rational thought'. Journal of Experimental Psychology: General, 127, 161-188.

Thompson, E., Chaiken, S., and Hazelwood, D. (1993), 'Need for cognition and desire for control as moderators of extrinsic reward effects: a person situation approach to the study of intrinsic motivation'. Journal of Personality and Social Psychology, 64, 987-99.

Tversky, A., and Kahneman, D. (1973), 'Availability: a heuristic for judging frequency and probability'. Cognitive Psychology, 5, 207-232.

Weatherson, B. (In Preparation), 'Defending interest-relative invariance'.

Weinberg, J. M., Alexander, J., Gonnerman, C., and Reuter, S. (2012). ‘Restrictionism and reflection: Challenge deflected, or simply redirected?’ *The Monist*, 95, 200-222.

Williamson, T. (2005), ‘Contextualism, subject-sensitive invariantism, and knowledge of knowledge’. *The Philosophical Quarterly*, 55, 213-235.

Wu, S. and Keysar, B. (2007), ‘The effect of culture on perspective taking’. *Psychological Science*, 18, 600–606.

¹ Authorship is equal; authors are listed in alphabetical order. We would like to thank Susan Birch, Steven Gross, Kareem Khalifa, Jennifer Nagel, and Jonathan Weinberg for helpful comments. We would also like to thank Ronald Mallon, Shaun Nichols, and the National Endowment for the Humanities for organizing and funding the 2012 Experimental Philosophy Institute where this project was conceived.

² We should note that the terms ‘project’ and ‘projection’ can be used in a number of different ways. In the Freudian tradition, paradigmatic cases of projection involve ascribing negative attributes to others without the projector being aware that he possesses these attributes (Gilovich, 1991, p. 113), a form of what David Holmes (1968) calls “similarity projection”. For others, projection has more to do with a particular account of third-person mindreading: it arises when the simulation processes of the attributor inappropriately fail to suppress a mental state of her own (e.g., Goldman, 2006, p. 165). In this paper, we are going to use the terms in a more general way; for us, projection will involve attributing one’s own mental state to others, whether or not

one is aware of having the state in question and however the psychology responsible for the attribution occurs.

³ There is a related controversy about whether knowledge attribution is sensitive to the personal costs of being wrong, and why. Subject-sensitive and interest-relative invariantists argue that knowledge is sensitive to what is at stake for the person whose epistemic situation is being described (Hawthorne, 2004; Stanley, 2005; Hawthorne and Stanley, 2008; Fantl and McGrath, 2010) and contextualists argue that knowledge is sensitive to what is at stake in the conversational context in which knowledge is being either attributed or denied (DeRose, 1999).

⁴ Hawthorne (2004) and Williamson (2005) provide a different psychological explanation, arguing that the relationship between our willingness to attribute knowledge and what possibilities have been made salient in a given conversational context can be explained in terms of well-known psychological bias called the availability heuristic. In addition to providing her own psychological account, Nagel (2010) argues that there are empirical and conceptual problems with an explanation grounded in the availability heuristic. We aren't interested in taking sides in this debate, and plan to focus only on the empirical merits of Nagel's alternative account. For more information about the availability heuristic, see Tversky and Kahneman (1973).

⁵ Epistemic egocentrism and the curse of knowledge are related to other well-known psychological biases, including the hindsight bias and the outcome bias. For additional discussion, see Hawkins and Hastie (1990) and Baron and Hershey (1988).

⁶ When discussing empirical results that suggest that perhaps folk knowledge attributions aren't sensitive to what is at stake for the person whose epistemic situation is being described in a given vignette, Brian Weatherson (in preparation) notes that invariantists don't think that subject stakes

always matter; instead, they simply think that subject stakes sometimes matter. It seems that a similar point could be made about salience, and would provide a way for invariantists to explain away the first wave of empirical results. Another way to explain away these results is by suggesting that some defeaters are properly ignored, a point made quite nicely by J.L. Austin (1990).

⁷ Five additional participants were excluded from the analysis: two for failing to answer the target question, and three for incorrectly answering a comprehension question. The exclusion criteria were decided before collecting data, and including their data doesn't affect the significance of the result.

⁸ There shouldn't be a significant difference with relatively small Ns. With really large Ns, however, matters may be different, if only because being told that John Doe is considering the possibility might result in a slightly stronger signal, compared to that resulting from automatic, largely unconscious mindreading processes. A similar comment holds for the next prediction.

⁹ Twenty-five subjects were excluded for failing the comprehension question, or failing to answer it at all, and thirty-five additional subjects were excluded for following basic directions. This is a large number of exclusions, however, the exclusion criteria were decided before the collection and analysis of data. Including these participants does not change the result ($F = 1.23$, $p = .301$).

¹⁰ Eleven participants were excluded from the analysis: one for failing to answer the target question; another for failing the comprehension question; and nine others for ignoring basic instructions. Exclusion criteria were decided before any analyses were run. Including the data does not change the significance of the result ($r = -.264$, $p = .007$).

¹¹ It is perhaps important to emphasize that the standard line on epistemic egocentrism is that we are unaware of the fact that the projection has occurred (Baron & Hershey, 1988; Camerer et al., 1989).

¹² A person's need for cognition is determined on the basis of her responses to a survey with eighteen self-report items like "I find satisfaction in deliberating hard and for long hours" or "I like tasks that require little thought once I've learned them." The eighteen-statement survey is a shortened version of the survey originally used by Cacioppo and Petty; however, studies have shown that the shortened version is just as accurate a measure as the longer version of a person's need for cognition. See Cacioppo et al. (1984).

¹³ Since person's need for cognition is supposed to represent her intrinsic motivation to engage in effortful thinking, she will find such an activity rewarding even if no external rewards are offered. In fact, Thompson et al. (1993) found that people with a high need for cognition will actually give less effort and care to cognitive tasks when they are presented with external rewards than when they are simply asked to engage in the cognitive tasks for their own sake.

¹⁴ Some caution is probably needed here. If Nagel's story is only a small part of the real story, and some other factors are also playing a significant role in generating these kinds of salience effects, then we might expect to find decreased tendencies to attribute knowledge in Nagel's more detailed case even when people aren't mistakenly projecting their own concerns onto others.

¹⁵ While the two-way ANOVA didn't reveal a main difference across conditions, it very nearly found a main difference across NFC groupings ($p = 0.059$). When we collapse the conditions, the

different NFC groupings seemed (somewhat) to be more or less permissive with their knowledge attributions: Low NFC ($\underline{M} = 4.82$), Mid NFC ($\underline{M} = 4.23$), and High NFC ($\underline{M} = 4.23$).