

Knowledge Ascriptions and the Psychological Consequences of Changing Stakes

It's late on Friday afternoon, and raining hard. Lo and her next-door neighbor Hi are thinking about going out to the bank, but wondering whether the trip could be postponed until tomorrow. Both of them can remember a recent Saturday visit to the bank, but neither of them has any further information relevant to the question of whether the bank will be open tomorrow. Nothing much is at stake for Lo—her banking errand could be done any time in the next week—but for Hi the question has burning practical importance. Hi knows he must deposit his paycheck before Monday, or he will default on his mortgage and lose his home. Does Lo know that the bank is open tomorrow? Does Hi?

If you are like most contemporary epistemologists, you find it easier to ascribe knowledge that the bank will be open Saturday to Lo than to Hi, despite the fact that the two subjects are relying on the same information in trying to settle that question. What's disputed are the reasons for such shifts. There are various rival explanations of why our natural inclinations to ascribe knowledge become more stringent or lax when subjects are described as, for example, having pressing practical interests in the proposition believed. Contextualists take such shifts to indicate that the referent of 'know' changes with the situation of the ascriber: when we read about Hi's predicament our epistemic standards become more demanding than they are when we read about Lo. In support of their focus on the ascriber, contextualists will point out that our sense of a shift is even stronger if each subject's situation is presented in a separate story, giving the ascriber a

better opportunity to adjust to the appropriate epistemic level. Meanwhile, advocates of subject-sensitive or interest-relative invariantism (IRI) focus on the subject to whom knowledge is ascribed. Rather than allowing that the referent of ‘knows’ can change with the conversational context of the ascriber, they argue that the relation picked out by ‘knows’ is always the same, but it is a relation in which the subject’s practical interests play a part unrecognized in traditional epistemology. Whether someone has knowledge is determined not only by truth-conducive features of belief formation such as her initial information and the reliability of the cognitive processes she employs, but also by such factors as her interests and the salience *to her* of various possibilities of error. Meanwhile, invariantists who maintain that only truth-conducive factors count (adhering to traditional ‘intellectualism’ rather than interest-relativity) typically attempt to dismiss these intuitions of shift, construing them as the products of psychological bias or failure to observe the distinction between knowing and knowing that one knows.¹ The aim of the present paper is to propose an intellectualist invariantist account that accommodates rather than dismisses our intuitive responses to these cases: I argue that our stringency with Hi and laxity with Lo is best explained by our natural recognition of differences in the truth-conducive features of high- and low-stakes belief formation.

Throughout the debate over shifting intuitions much attention has been paid to the question of whether there is empirical linguistic evidence for context sensitivity in knowledge claims. The lessons to be learned from examining the lexical semantics of ‘know’ are subtle and arguably still inconclusive [DeRose 2005; Stanley 2005; Ludlow 2005]. Meanwhile, less attention has

¹ This Hi/Lo case is a version of the bank case originally developed by Keith DeRose in support of contextualism, and adapted to support IRI by Jason Stanley. Contextualism has been defended in Lewis 1996; Cohen 1999; DeRose 1992; IRI in Stanley 2005, and, more guardedly, Hawthorne 2004; and strict invariantism in Williamson 2005 and elsewhere. Other ways of handling comparable epistemic shift intuitions include contrastivism [Morton and Karjalainen 2003; Schaffer 2005] and relativism [Richard 2004; MacFarlane 2005].

been paid to findings in empirical psychology that might bear on the question of how to interpret the knowledge ascription data.² Given the difficulty of interpreting the problematic ascription patterns, and the suggestion that some intuitive ascriptions might be due to heuristics or biases, there is some value in examining the psychological mechanisms behind knowledge ascription more closely. In what follows I examine some psychological mechanisms relevant to differences in belief formation under high- and low-stakes conditions. The differences are pronounced enough to bar us from assuming that subjects who have the same initial information but different stakes are naturally perceived to enjoy belief formation that is the same as far as truth-conducive factors are concerned.

Focusing on some specific examples from the recent epistemological debate, I argue that our sense of a shift in epistemic quality in the DeRose/Stanley bank cases can be explained by a psychological effect often causally related to, but conceptually distinct from, perceived stakes. The effect is known as ‘need-for-closure’, where ‘closure’ marks the switch from the formation to the possession of a belief. Prior to closure we are trying to figure out what to believe, searching for information and weighing various alternatives; after closure we have the subjective sense of a solid result. Our level of need-for-closure sets the point at which we can settle on an answer to a given question; we can think of it as a thermostat setting the temperature at which fluid evidence assessment will freeze into solid belief. One can manipulate how high the

² John Hawthorne discusses some work on biased judgments of risk as one possible explanation of a tendency to ‘overproject’ one’s sceptical sentiments in assessing the beliefs of others [Hawthorne 2004: 164]; he suggests our mere contemplation of problematic counter-possibilities may raise our estimation of their likelihood [cf. Williamson 2005]. Stewart Cohen, Keith DeRose and Jonathan Schaffer argue that there is no clear advantage for IRI over contextualism in the psychological work cited by Hawthorne [Cohen 2005, DeRose 2005; Schaffer 2006]. Cohen points out that the empirical data Hawthorne cites does not quite support the suggestion that mentioning a risk always elevates one’s estimation of its likelihood, and there is further evidence -- e.g. [Sherman et al., 1985] -- that imagining a risk does not have a uniformly positive impact on its perceived likelihood. While the focus of the current paper will be strictly on practical interests, epistemologists also need to do more to examine what happens to us psychologically when possibilities of error are mentioned.

thermostat is set by changing the subject's perceived interests (say, penalizing him for inaccuracy), or, equally well, by changing environmental conditions like time pressure and background noise. Because subjects with different levels of need-for-closure can differ in whether they have or lack knowledge for the straightforward reason that they have or lack a solid belief, we do not need to resort to non-traditional epistemology to explain the shifts produced by practical interests in cases like the Hi/Lo one at the outset of the paper. To be sure, these cases can be amended to stipulate that the subjects with different interests and the same initial information have the same level of confidence in the target belief, but I'll argue that the most psychologically plausible way of understanding the amended cases will produce contrasting epistemic intuitions (Lo knows, Hi doesn't know) only if the high stakes subject is perceived as having elevated need-for-closure, as being under a condition like time pressure or distraction. Because haste and distraction lower accuracy, the intellectualist invariantist can explain these amended cases as well: it's not surprising that apparently hasty Hi looks bad beside his unrushed counterpart, even if both of them have the same initial information and are giving the same answer. If cases of contrasting epistemic intuitions involving contrasting practical interests can always be interpreted as cases of compromised accuracy, then such cases cannot provide evidence for contextualism or IRI over intellectualist invariantism.

1. High and Low-stakes thinkers

Here's an unsurprising psychological fact: when asked to solve the same problem, high-stakes subjects tend to try harder than their low-stakes counterparts. Psychologists have used a wide array of methods to manipulate their subjects' perceived stakes, ranging from offering financial rewards for accuracy to telling subjects that that they will be asked to justify their choices in

front of others, or that their performance on a given task is correlated with the possession of desirable qualities. A variety of high-stakes effects have been uncovered, and the specific ways in which stakes are raised can make a difference to the effects produced, but some general results seem to apply across the board. When it is possible to do so, high-stakes subjects will search for more information before making up their minds, and process this information more thoroughly, for example by selecting more complex decision strategies, such as assigning different weights to the more and less important attributes of the products they are comparing rather than just weighing all attributes equally.³ In general, high-stakes subjects think more systematically and less heuristically, relying more on deliberate and controlled cognition and less on first impressions and automatic responses [Kunda 1990; Lerner and Tetlock 1999]. Many cognitive biases – a recent survey article on accountability counts sixteen – are known to be attenuated when subjects take themselves to be shifted into a higher-stakes condition [Lerner and Tetlock 1999]. The flaws diminished by perceived higher stakes include the primacy bias, in which we are oversensitive to the order in which information appears, the overattribution bias, in which we overestimate dispositional causes and undervalue situational causes, and the sunk-cost error, in which we adhere irrationally to decisions made in the past.

Although there is a general correlation between perceived high stakes and more accurate cognition, it seems that what really increases under the perception of high stakes is cognitive effort rather than accuracy itself. Several minor biases—notably those whose correction requires knowledge of formal statistical rules—are unaffected by accuracy incentives, and some biases

³ On increased searching for information, see e.g. [Sanitioso and Kunda 1991; Huneke et al. 2004]; on more complex cognitive strategies to reach a judgment, both when cued and spontaneously, see e.g. [McAllister, Mitchell, and Beach 1979; Van Hiel and Mervielde 2003].

are differentially affected depending on the degree to which they involve effortful thought.⁴ In certain unusual situations, the increased effort of high stakes subjects can even result in lowered accuracy. On the other side of the ledger that found sixteen biases to be diminished by higher stakes, five biases came out as amplified when stakes are high, most notably the dilution effect, in which high stakes subjects assign weight to all information provided, whether or not it is actually relevant to the task. In one study, high-stakes subjects asked to estimate a student's GPA were disproportionately guilty of being influenced not only by genuinely diagnostic information (e.g. hours spent studying) but also by such extraneous details as the number of houseplants the student was said to own [Tetlock and Boettger 1989]. However, this last impairment of judgment is arguably a response to the way the conversational norm of relation applies to questions posed in the laboratory, where subjects might have expected all information supplied to have value. The high-stakes dilution effect increases when the conversational norm is explicitly primed (subjects are told that 'all information has been carefully selected') and vanishes when the norm is explicitly cancelled ('some information provided may be irrelevant') [Tetlock, Lerner, and Boettger 1996]. In the real world, where we have no standing presumption of the relevance of all information encountered, it is not clear that high stakes subjects tend to suffer dilution.

Especially where real world high-stakes subjects are at liberty to indulge their tendencies to collect more information prior to judgment, their greater effort should typically be expected to

⁴ The numerical anchoring bias, for example, is insensitive to financial incentives for accuracy when the anchor is supplied by the experimenter, but attenuated by incentives when the anchor is generated by the subject's own efforts. In anchoring, judgments are biased in the direction of a cue: subjects first asked whether Mount Everest is more or less than 2,000 feet tall went on to give a median estimation of 8,000 feet as its height; subjects who were instead asked whether it is more or less than 45,000 feet tall gave a subsequent median estimation of 42,550 feet.[Jacowitz 1995] The effect persists even when subjects are forewarned about the effect, and when it is made vivid to subjects that the anchor is randomly selected, for example on the basis of the spin of a wheel of fortune [Tversky and Kahneman 1974; Chapman and Johnson 2002]. The anchor does not have to be provided by the experimenter, however: a question like, 'In what year did the second European explorer land in the West Indies?' prompts the subject to think of 1492 on her own, and then adjust upwards from there. Epley and Gilovich have shown that financial incentives and forewarnings reduce anchoring effects only when the anchor is generated by the subject [Epley and Gilovich 2005]. They argue that the more belief formation is depends on effortful thinking rather than automatic and unconscious processes, the more sensitive it will be to perceived incentives.

pay off in greater accuracy. Now, given that high-stakes subjects tend to be more accurate, it may seem odd that we are naturally *less* willing to ascribe knowledge to them than to their less accurate low-stakes counterparts, but high stakes subjects pay a price for their higher accuracy – slower and less confident belief. Where a Low Stakes subject would have settled already, the High Stakes subject is still assessing his evidence and searching for more.

Before turning to the application of these contrasts to epistemological cases, there is one more effect of the increased cognitive effort of high-stakes subjects that deserves special mention: these subjects have greater subjective awareness of their processes of perception, memory and judgment. The Source Monitoring Framework (SMF) developed by Marcia Johnson is one standard account of the phenomenon [Johnson, Hashtroudi, and Lindsay 1993; Johnson and Raye 2000]. Given a cue or problem ('Is the bank open this coming Saturday?') which activates certain information (say, thoughts of a recent weekend trip to the bank), all subjects pass through a stage of evaluating this information in order to reach the judgment that the trip really did happen in the quite recent (as opposed to immediate or distant) past, and that the source of one's present thought is indeed memory of perceptual experience, rather than say, memory of some past visualization or contemplation of a possible trip that did not end up occurring. According to the SMF, the initially activated information carries no distinct tag whose function is to indicate that its origin was in veridical 13-day-old experience rather than memory of an event from an earlier or later time, or from some rival source such as an episode of imagination. Rather, studies of source monitoring and its susceptibility to manipulation indicate that subjects determine source at the time of recall by assessing a range of features of the content and quality of the activated information, in a manner which differs with certain circumstantial factors. Importantly

for our present purposes, one such factor is the subject's perceived interests: low-stakes subjects ordinarily perform source attributions automatically and heuristically, guided by surface characteristics of the activated information, such as degree of visual detail. High-stakes subjects supplement the heuristic processing with more deliberate and effortful reasoning, for example by checking the plausibility of what is recalled against other beliefs and by attempting to develop a more precise representation of the temporal context of the event. Stakes influence not only whether the source decision is heuristic or systematic, but also the amount of information required in either kind of decision. In Johnson's words, 'Depending on our motives, the types of information we expect or require to make a source decision will vary, whether we are engaging in a heuristic process or systematic process or both.' [Johnson and Raye, 2000: 40] The very same information might in low stakes contexts suffice to produce confidence that what is recalled is veridical while failing to do so when stakes are high.

That shifting from low to high stakes has an impact on belief formation does not automatically imply that our intuitive knowledge ascriptions tend to reflect this impact. It is possible that we tend to be blind to these changes in belief formation, either in first- or the third-person cases; it is also possible that we are aware of these sorts of changes but that our shifting epistemic intuitions are driven by other forces. Still, where it explains an epistemic shift to say that we expect increased cognitive effort from a high-stakes subject, the hypothesis that such an expectation is driving our sense of shift is one that deserves to be kept in mind. Given that it's a reasonable and accurate expectation, there is no special difficulty in explaining how it would come to be operative. Advocates of contextualism and IRI have an additional reason to hope that our intuitive knowledge ascriptions actually reflect the high-stakes/high-effort correlation: to the

extent that one's epistemology is driven by a desire to accommodate intuitive knowledge ascriptions, the more these ascriptions capture significant differences in actual belief formation, the better.

It should also be emphasized that the existence of a correlation between stakes and effort is not itself any bar to accepting a theory like contextualism or IRI. Contextualists can point out that the current setting of one's own thermostat should have an impact on whether one views others' beliefs as adequate or half-baked; they can also emphasize that some of the core cases for their position concern not practical interests, but possibilities of error (an area untouched in the present paper). Meanwhile, advocates of IRI who take knowledge as the goal of inquiry should be pleased to note that we work harder to attain that goal when stakes are high. However, I'll argue in what follows that if certain stake-related psychological effects actually explain the epistemic shifts in the practical interest cases used to support contextualism and IRI, these cases do not have the evidential value against intellectual invariantism that they have been claimed to have. If we can't assume that subjects with the same initial information and different stakes are naturally perceived to enjoy belief formation that is the same as far as truth-conducive factors are concerned, cases that rest on that problematic assumption need to be refashioned. After arguing that existing cases do count on the problematic assumption, I'll argue that the requisite refashioning will not help contextualism or IRI: when robust epistemic intuitions shift in response to varying practical interests, the intellectualist invariantist will always have a psychologically plausible story about why this is so.

2. Some epistemological test cases

The cases that are said to provide the clearest evidence for contextualism and IRI are cases involving subjects who are reasonable, aware of their practical interests, and given the same initial information upon which to form a belief, with a ‘pressing and very practical concern’ in the High case, and no corresponding practical concern in the Low case.⁵ Here is such a case, originally developed by an advocate of contextualism, and retold by an advocate of IRI:

Low Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it isn’t very important that their paychecks are deposited right away, Hannah says, ‘I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning.’

High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But as Sarah points out, banks do change their hours. Hannah says, ‘I guess you’re right. I don’t know that the bank will be open tomorrow.’ [Stanley 2005: 3-4]

Jason Stanley finds that when given the information that the bank will be open, we intuitively judge that Hannah is right both in claiming knowledge in Low and in denying it in High; he contends that this shift occurs because interests are seen to figure as a factor in knowledge.

Arguing against IRI, Jonathan Schaffer notes that the cases presented differ not only in stakes but in other ways: for example, only in High is there mention of the possibility of error. Schaffer goes on to present stripped-down variants of High and Low with no talk of error. For his versions Schaffer tends to judge knowledge can be attributed in both Low and High, and notes that

⁵ The quoted phrase is DeRose’s, who joins advocates of IRI in finding that these cases elicit ‘stronger and more stable’ intuitions than cases not involving practical interests, or cases in which subjects or other parties to the conversation are unreasonable or mistaken about their interests [DeRose, 2005].

whether or not one judges the cases uniformly, ‘the intuitions are nowhere near as clear’ [Schaffer 2006: 89].

Advocates of IRI can grant that less detailed cases might trigger weaker intuitions, and perhaps plead that the expressed concern with the possibility of error in High Stakes serves to draw attention to Hannah’s practical interests.⁶ Indeed, it’s unclear that making the cases parallel must favour contextualism: another way to do this would be to add a reminder by Sarah about the possibility of a change in hours to Low, and with such a change I experience the original IRI intuitions.⁷ But however things stand for IRI, Schaffer is certainly right that more work is required to understand why the original cases generate strong intuitions of shift.

Kent Bach argues that what really changes between Low and High is the subject’s threshold for confident belief in the proposition that the bank is open on Saturday: perhaps in High her true belief is insufficiently confident to count as knowledge [2005].⁸ It is easy to read Hannah as having less confidence in High; indeed, one could make the stronger claim that the cases are really contrasted on whether Hannah has at the end of the case yet attained a firm belief on the target proposition at all. Hannah affirms knowledge of her conclusion in Low but in High asserts

⁶ Advocates of IRI could also object to certain features of Schaffer’s version. In both cases Schaffer pops the question of whether the subject knows immediately after a sentence declaring, ‘He is right – the bank will be open.’ Because this affirmation follows a brief description of the subject’s inferring that the bank will be open from his memory of last week’s opening, it’s hard not to pick up a hint that the subject’s *inference* is right, to read this as an affirmation that the thinking in this case is adequate. Schaffer’s version also reduces the temptation to shift by making the evidence fresher and the odds of a change in hours lower (6 days have elapsed rather than 13).

⁷ With such a change, it appears to me that Low Stakes Sarah is being neurotic in worrying about error, not that Low Stakes Hannah lacks knowledge. Schaffer also develops a pair of cases mentioning the possibility of error, and contends that the shift will disappear into a uniform denial of knowledge, but my intuitions on these cases are very weak. These cases also differ from Stanley’s in inserting the suggestion of error in the voice of the narrator rather than the conversational partner of the subject, which makes a difference to the expected subjective confidence of the subject (a difference to be discussed shortly).

⁸ Jon Kvanvig is also credited as having presented a similar suggestion on his blog *Certain Doubts*. See [Stanley 2005: 6].

only her evidence for it, as if she is in the process of figuring out what to believe about the coming Saturday's banking hours. If she is, the contrast between High and Low is not a contrast between knowledge and true belief, but a contrast between knowledge and a state in which the assessment of evidence has not yet solidified into belief. The truth of the subjunctive conditional that High-Stakes Hannah would, if she were pressed, express a belief that the bank will be open rather than closed does not indicate that she is represented in the existing scenario, where such pressure is absent, as having arrived at that judgment at the moment we are invited to assess her.⁹

Whether one thinks of High-Stakes Hannah as having a lower-confidence belief or a state of cognition that precedes arrival at settled belief, it's psychologically realistic to read her as needing more evidence either to make up her mind at all or to attain the same level of subjective confidence in High Stakes as she would have in Low. As a high-stakes/high effort subject, she should be expected to collect more information prior to settling on a fixed belief. But we need a sharper view of the relation between interests and the way we settle on a fixed belief to figure out what is going on with our inclinations to ascribe knowledge here. This is an area that has been investigated in detail by psychologist Arie Kruglanski, and I think a concept of Kruglanski's – the concept of *need-for-closure* – is useful in explaining the shift.

Kruglanski introduces the term 'closure' as a name for arrival at a settled belief: in his words, closure is 'the juncture at which a belief crystallizes and turns from hesitant conjecture to a subjectively firm "fact"' [Kruglanski and Webster 1996: 266]. Achieving closure or judgmental

⁹ As further evidence that Stanley's cases might have to do with the difference between the presence and absence of settled belief, it's worth noting that for Schaffer's stripped-down versions of High and Low, which more readily elicit the reaction that the subject knows in both cases, the subject is explicitly said to draw the conclusion that the bank will be open on the coming Saturday.

commitment on a question puts an end to the experience of ambiguity and delivers the sense of having a firm answer. The opposite of closure is openness or judgmental non-commitment, in which we are able to continue juggling alternative possibilities, perhaps lingering in ambiguity or confusion. A central interest of Kruglanski's is the extent to which motivational factors influence the attainment or avoidance of closure.

Kruglanski starts by drawing a distinction between motivation or need for specific closure and need for non-specific closure. Need for specific closure is what philosophers tend to think of in connection with the impact of motivation on the formation of belief: this is the motivational force that biases us in wishful thinking, and, more generally, drives us whenever arriving at some particular answer (or type of answer) is antecedently considered very desirable. Less philosophical attention has been paid to motivation for non-specific closure, which is the drive to have some answer or other, where what is antecedently desired is just the state of settled belief itself [Kruglanski and Webster 1996; Kruglanski 1989]. Need for specific closure drives the patient scanning the internet in the hopes of finding out that his condition is curable; need for non-specific closure drives the pathologist who is trying to complete a report and simply wants to know whether some anonymous patient's tumour is malignant or benign, where neither outcome is antecedently preferred by the pathologist. Unsurprisingly, motivation for specific closure has a significant impact on belief formation, starting with various wishful thinking effects. (For a comprehensive survey, see [Kunda 1990].) But it is motivation for non-specific closure that will occupy us in what follows: what Kruglanski says about this type of motivation has particular value for cases like the Bank case.¹⁰

¹⁰ From here on I'll be using 'closure' to mean non-specific closure. Note that motivation for specific closure and motivation for non-specific closure are not exclusive; it's possible to increase both forms of motivation

Motivation for non-specific closure is something that Kruglanski envisages as ranging in a continuum from a high level of need-for-closure (where action must be taken immediately, further ambiguity is intolerable, or the costs of further cognition are high), through a neutral condition, and down to low need-for-closure, also called ‘need to avoid closure’ (where the subject has what Kruglanski dubs a ‘fear of invalidity’, a strong desire to avoid freezing on an inaccurate result, or even just desires to continue reflecting and avoiding the attainment of any fixed result at all). This type of motivation can be increased or decreased in a number of ways. High need-for-closure can be induced by increasing the costs of continuing in ambiguity, for example by putting subjects under perceived time pressure, fatigue, raising the level of background noise, or making the problem-solving task seem tedious. Low need-for-closure can be induced by making subjects strongly averse to inaccurate or premature judgment, as in the High Stakes bank case. The reverse orientation of Kruglanski’s high/low contrast makes the terminology confusing at first, so an explicit review may be helpful. Other things being equal, high stakes subjects experience low need-for-closure, and low stakes subjects experience neutral (neither high nor low) need-for-closure. Cases in which subjects are described as rushed or distracted can count as high need-for-closure cases, or at least as cases in which subjects have higher need-for-closure than their unrushed counterparts.

Conveniently, self-esteem concerns can be used for manipulations in either direction: in one study, non-control subjects were told that mental concentration and intelligence were correlated

simultaneously, for example, by changing both incentives to reach a certain outcome and incentives for accuracy, or by adjusting time pressure or environmental conditions. Perhaps surprisingly, motivation for specific and non-specific closure can function orthogonally. Unfortunately, the Bank and Airport cases involve both kinds of motivation; for a more fine-grained understanding of intuitive knowledge ascription one would also need to develop sets of cases that involve just one sort of closure motivation at a time.

with either clear-cut opinions (to induce high need-for-closure) or accuracy (to induce low need-for-closure) on a specific task [Maysseless and Kruglanski 1987]. The task involved the recognition of common household items from enlarged photos of parts of these objects, shot from odd angles. Subjects were encouraged to list as many hypotheses as possible about what the photographed objects might be, and then indicate which hypothesis was most likely right. Given equal time and the impression that intelligence was being assessed, all subjects might be expected to generate similarly impressive lists of hypotheses. In fact, subjects under a high need-for-closure condition generate markedly fewer hypotheses than controls, who in turn generate fewer than those in the low need-for-closure condition. Kruglanski notes that this difference in hypothesis generation should itself produce higher subjective confidence on the part of the high need-for-closure subjects, given the Kelley discounting principle, according to which subjective confidence in a given hypothesis will tend to drop simply as a function of number of alternative hypotheses available [Kelley 1971].

A more direct test of the relation between subjective confidence and need-for-closure was carried out in a separate experiment, in which participants in high and low need-for-closure were asked to identify digits flashed very briefly on a screen. All subjects were given control of the projector and allowed to repeat the stimulation as often as they wished. Subjects in the high need-for-closure condition chose to curtail their collection of evidence before controls, who in turn collected less evidence than those in the low need-for-closure condition; furthermore, both confidence in initial hypotheses and confidence shifts on increasing evidence were reported as highest in the high need-for-closure condition, intermediate in the control and lowest in the low need-for-closure condition [Maysseless and Kruglanski 1987].

The fact that elevated need-for-closure generates greater confidence on less evidence is part of what Kruglanski calls the ‘unfounded confidence paradox’, which manifests itself in a number of ways. Higher subjective confidence can be produced by putting participants under perceived time pressure, subjecting them to distracting background noise, or making them think of their task as tedious. When cognition seems costly, we want to get it over with; we think up fewer hypotheses and feel better about the one chosen. Subjects who solved the same problems without sensing time pressure, in silence, or while thinking of their task as neutral or pleasant felt less subjective confidence despite giving more accurate answers (less influenced, for example, by the primacy bias, in which we are oversensitive to the order in which information appears).¹¹

Kruglanski dubs the compromised thinking under elevated need-for-closure ‘seizing and freezing’: in this condition we pounce on information to form a belief and then become resistant to evidence contrary to the conclusion first attained. By contrast, in low need-for-closure we delay judgment and explore evidence more thoroughly (and typically with greater accuracy).

High Stakes Hannah is a low need-for-closure subject: she is strongly averse to making a mistake about the banking hours, and willing to entertain hypotheses that would make her initial evidence inconclusive. Her Low Stakes counterpart is in a neutral need-for-closure condition: Low stakes Hannah is not described as being under anything like the pressure for immediate decision characteristic of high-need-for-closure subjects, nor is she strongly averse to making the wrong call about Saturday banking – the scenario suggests that getting it wrong would be a mild inconvenience at worst. If our intuitions about the shift in Stanley’s cases are driven by the

¹¹ Time pressure: [Kruglanski and Webster 1991]; background noise: [Kruglanski, Webster, and Klem 1993]; affective perception of task: [Webster, 1993].

contrast between low and neutral need-for-closure, then something like the Bach objection will be right: we ascribe knowledge in Low Stakes and deny it in High because we naturally attribute higher and lower confidence belief to the contrasted subjects, or confident belief and a state of evidence assessment that precedes fixed belief. Because of the perceived need-for-closure differences between the subjects, we don't expect the same information to produce the same level of belief in High and Low Stakes. It is true that a difference in stakes causes this difference in our knowledge ascriptions, but it does so in a way that should leave the traditional intellectualist epistemologist untroubled: whether or how strongly a subject believes something is a squarely traditional component of knowledge.

If these cases depend on the level need-for-closure caused by the subjects' interests, rather than the interests themselves, we should also expect our intuitions to shift again if need-for-closure is independently manipulated, for example by time pressure. Conveniently, Schaffer has developed a version of the bank cases altered in just this way. In his *Low and Slow*, Sam, not wanting to stand in line on Friday, believes that the bank will be open on Saturday on the basis of remembering last week's opening. He has no pressing need to deposit the check Saturday. Thanks to car failure he finds himself right in front of the bank on Friday, and chooses to spend an hour resting in his car right there rather than taking the easy opportunity to confirm the banking hours. Meanwhile, in *High and Fast*, Sam has a similar desire not to stand in line, and similar evidence, but a pressing need to deposit the check before Monday. As he is about to double-check the bank hours he remembers that he has a very small window of opportunity to buy his wife a present ('his whole relationship is at stake'), and must make a split-second decision either to check the hours or buy the present. He passes up his chance to check the hours

in order to buy the present. Schaffer tends to intuit (contrary to what he thinks IRI has to say) that Sam lacks knowledge that the bank will be open Saturday in *Low and Slow* but has this knowledge in *High and Fast* [Schaffer 2006].

I think *Low and Slow* is not terribly problematic for IRI. The suggestion that a subject is deliberately avoiding available evidence is enough to generate a concern that his belief is ill-formed on any theory, perhaps driven by a desire not to stand in the long Friday line. Even if IRI is right that low stakes make knowledge ascriptions less stringent, they can still be stringent enough to censure conspicuous wishful thinking. Our knowledge-ascribing intuitions in *High and Fast* are more problematic for IRI. Sam is forming his belief about the bank opening on the same evidence as the original High Stakes Hannah, and has (roughly) the same practical concern in the matter; if representing him as under time pressure is enough to switch us over to the intuition that he does know the bank will be open, then it appears practical interests themselves are not what count in this case. If time pressure, a condition raising need-for-closure, is all we need to diminish the stringency of a High-Stakes knowledge ascription, then we have reason to believe that the original stringency was due to the fact that such cases ordinarily involve a low need-for-closure condition (a condition in which the subject would ordinarily seek out more evidence and assess it more thoroughly before settling upon a belief). High stakes situations ordinarily cause low need-for-closure, so it is easy to think that high stakes themselves are what matters, but when these conditions are pulled apart we see that what really makes the difference is not the stakes but the perceived level of need-for-closure. So Stanley's bank cases do not show that interests themselves make a difference, with truth-conducive factors and subjective confidence held fixed.

3. Stipulating confident belief

It is one thing to argue that interests don't underpin one example, and another thing to argue that they don't ever count the way they were supposed to. Perhaps advocates of IRI and contextualism just need better examples; perhaps the impact of stakes on knowledge or knowledge ascriptions is real but rather uncommon, and it takes more unusual situations to bring it out.¹² Even if Stanley's Bank cases really contrast high and lower confidence belief, or confident belief with a state of evidence assessment, we can develop psychologically plausible cases in which we stipulate that Lo and Hi form the same belief with the same degree of confidence, and in which it seems to us that Lo knows and Hi does not. In fact, the paradox of unfounded confidence gives us a simple recipe for doing this. Consider a variant of Schaffer's High and Fast:

High Stakes, high need-for-closure Mack. Mack is an oncologist. It's late in the day, he is behind in his work, and needs to clear one more case before he leaves. He reviews some highly complex evidence, pronounces a diagnosis, and rushes out the door to go buy his husband a present before the shops close. The evidence was of a sort that doctors sometimes misread. Someone's life hangs in the balance.

Low Stakes, low need-for-closure Jack. Jack is Mack's twin brother, and has had identical medical training. He is calmly taking a practice test with no consequences for himself or others. He reviews some highly complex evidence of a sort that the doctors sometimes misread (say, the actual patient file from the Mack case). He is enjoying the experience of writing the practice test (a nice break from seeing patients), the room he is writing the test in is quiet, and he is not under any perceived time pressure. Jack eventually arrives at a diagnosis, and when he does so he feels the same level of confidence as Mack had.

It is easier to ascribe knowledge to Jack than to Mack, despite their equal confidence and initial information. The advocate of IRI could argue that reason why our epistemic intuitions shift here

¹² Of course IRI does not insist that every difference in stakes makes a difference in knowledge ascription: it's consistent with Stanley's formulation that such shifts could be quite rare, and involve only fairly complex judgments.

is that given his stakes in making a diagnosis for a real patient, the oncologist needs to think harder in High than in the practice context of Low.

But the most obvious explanation of why we find the story about Mack psychologically plausible also explains why such cases will not ultimately help IRI. The case is most easily read as turning on the unfounded confidence effect: where a high-stakes subject is under time pressure or distraction (the scenario suggests both), these factors work to eliminate his high-stakes caution, the caution that would ordinarily have the result that his belief would have less than the stipulated level of confidence. However, given that haste and distraction typically lower the accuracy of judgment, there is a perfectly traditional explanation of why we become more stringent in our knowledge ascriptions moving from Low to High. For example, the hasty doctor would be more susceptible to the primacy bias, so presenting the evidence of the case to him in another order might have resulted in a different diagnosis: his correct judgment depended on his luck that the lab reports were stacked the way that they were.

Returning to the bank cases, we could stipulate that High-Stakes Hannah has not only the same initial information but also the same final confidence level as Low-Stakes Hannah. Given the differences in confidence ordinarily produced by differences in stakes, it's odd that High-Stakes Hannah is so confident, and one has to wonder how this peculiar confidence is understood by the reader of the case. One natural hypothesis that could explain the reader's tendency to ascribe a failure of knowledge to Confident High-Stakes Hannah is that the reader will naturally perceive her as having unfounded confidence of the type produced in Kruglanski's paradox of unfounded confidence. Because High Stakes subjects ordinarily experience low need-for-closure, it's only

psychologically plausible that High Stakes Hannah could have high confidence of the sort stipulated on her slim evidence if she is compromised in her accuracy, for example by thinking hastily or in a way biased by wishful thinking. If we imagine High Stakes Hannah asserting *I'm just sure it will be open* after she's warned of the chance the bank has changed its hours, what conclusions do we naturally draw about the character of her cognition? That her cognition would be seen to fall short of what is needed for knowledge because of truth-conducive factors is a hypothesis the advocate of IRI needs to eliminate before claiming evidential support from the case, but it's not clear how he will do so within the bounds of psychological plausibility.

Ordinarily, high- and low-stakes subjects think differently about problems complex enough to generate contrasting epistemic intuitions: we can make high-stakes subjects think like low-stakes subjects on such problems, but only by putting them under conditions where they think less accurately than their low-stakes counterparts. To the extent that our knowledge ascriptions are sensitive to stakes through being sensitive to lowered accuracy, they cannot be used by the advocate of IRI to establish the most provocative part of his thesis, namely that knowledge itself involves certain factors that are irreducibly practical rather than truth-conducive. Need-for-closure effects connect practical and truth-conducive factors, in a way that the advocate of IRI needs to watch.

One might hope that these accuracy-compromising effects could simply be taken out of the equation by considering cases in which a subject is described as unaware of her strong practical interests in a given question. Here, the subject's own perception of the situation is evidently not going to trigger the ordinary cognitive caution induced by perceived high stakes. These cases are

puzzling, and generate weaker and more disputed intuitions than the transparent cases discussed so far.¹³ For example, I don't see the pathologist as needing to do any more, cognitively, to know a diagnosis when the anonymously numbered tissue sample is from a loved one. One can try to cultivate IRI-type responses to such cases by imagining what extra steps the pathologist might take if only he knew his interests—and if the case is told with the right dramatic flair, it might become irresistible to do this—but this kind of counterfactual reasoning makes us assess ignorant high stakes cases as if they were perceived high stakes cases, and in doing so makes need-for-closure expectations relevant again. When subjects don't exhibit the caution we intuitively (even if wrongly) expect of them, they appear to be over-hasty. The standing temptation to assess ignorant high stakes cases as if they were perceived high stakes cases can be explained in terms of the hindsight bias: it is psychologically very difficult to suppress our knowledge of the subject's stakes in evaluating her reasoning.¹⁴ And as further evidence that such cases should be treated with caution, IRI does not yield any nice way of handling Ignorant Low Stakes cases, in which subjects falsely believe themselves to be in High Stakes but for some reason engage in cognition typical of Low Stakes. It is not hard to think of such cases in which the Ignorant Low Stakes subject will appear *not* to have knowledge despite having reasoned in a way appropriate to the actual stakes: again, we are (reasonably) inclined to be harsh about any case in which the subject lacks the ordinary cognitive caution characteristic of his perceived High Stakes condition, given that when such caution is undercut, knowledge-compromising forces like haste and distraction are typically at work.

¹³ For example, in his [2006] Schaffer gives a number of Ignorant High Stakes cases with Low Stakes counterparts which elicit no intuitions of the epistemic shift IRI predicts.

¹⁴ The hindsight bias impairs us from setting aside what we know while evaluating the perspective of another who should be taken to lack this knowledge. (It is asymmetrical – we have no difficulty 'subtracting ignorance' and reasoning about the perspective of others taken to be better informed than we are.) The bias cannot be cancelled by forewarnings or financial incentives for accuracy. For a review of the impact of the hindsight bias on knowledge ascriptions, see [Nickerson 1999]. It has been argued that hindsight, or 'the curse of knowledge', is the central psychological limitation on mental state reasoning [Birch and Bloom 2004].

What advocates of IRI and contextualism really need is a case in which we have a natural and robust intuitive shift from ascribing to denying knowledge across a pair of reasonable subjects differing only in stakes, a pair for whom we have stipulated not only that the same initial information is provided and a final belief is held with the same degree of confidence, but also that there is no difference in the processing of the information, no haste or distraction on the part of the subject whose belief fails to count as knowledge. I have not been able to construct any such cases; the inclination to ascribe knowledge to the low-stakes subject but not the high-stakes one does not persist when their cognitive situations are explained in full detail.

There are still many cases of apparent epistemic shift untouched by the argument of this paper, including contextualist cases involving the selective mention of possibilities of error, and cases in which the judgment of a single subject is assessed from the perspective of observers with different stakes or standards. However these further cases come to be explained, continuing attention to the psychology of belief formation and knowledge ascription should help us understand why certain patterns of knowledge ascription seem natural to us. In turn, understanding why certain patterns of knowledge ascription seem natural to us should help us gauge the distance between these everyday appearances of knowledge and knowledge itself.¹⁵

¹⁵ For comments on earlier drafts of this paper I am grateful to Kent Bach, Cheryl Misak, Diana Raffman and Sergio Tenenbaum. A version of this paper was presented at the Central Division APA in April of 2007; I owe thanks to members of the audience for their feedback and special thanks to Jason Stanley, who served as commentator, both for his astute comments on that occasion and for a thought-provoking discussion of these issues a year earlier.

References

- Bach, Kent 2005. The Emperor's New 'Knows', in *Contextualism in Philosophy*, ed. G. Peter and G. Preyer. New York: Oxford University Press.
- Birch, Susan and Paul Bloom 2004. Understanding Children's and Adults' Limitations in Mental State Reasoning, *Trends in Cognitive Sciences* 8: 255-260.
- Chapman, Gretchen B. and Eric J. Johnson 2002. Incorporating the irrelevant: Anchors in judgments of belief and value, in *Heuristics and Biases: The Psychology of Intuitive Judgment*, ed. T. Gilovich, D. Griffin and D. Kahneman, New York: Cambridge University Press: 120-138.
- Cohen, Stewart 2005. Knowledge, Speaker and Subject, *Philosophical Quarterly* 55: 199-212.
- Cohen, Stewart 1999. Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives* 13: 57-89.
- DeRose, Keith 2005. The Ordinary Language Basis for Contextualism and the New Invariantism, *Philosophical Quarterly* 55: 172-198.
- DeRose, Keith 2004. Single Scoreboard Semantics, *Philosophical Studies* 119:1-21.
- DeRose, Keith 1992. Contextualism and Knowledge Attributions, *Philosophy and Phenomenological Research* 52: 913-929.
- Epley, Nicholas and Thomas Gilovich 2005. When Effortful Thinking Influences Judgmental Anchoring: Differential Effects of Forewarning and Incentives on Self-Generated and Externally Provided Anchors, *Journal of Behavioral Decision Making* 18: 199-212.
- Hawthorne, John 2004. *Knowledge and Lotteries*, New York: Oxford University Press.
- Huneke, Mary, Catherine Cole and Irwin Levin, 2004. How Varying Levels of Knowledge and Motivation Affect Search and Confidence during Consideration and Choice, *Marketing Letters* 15: 67-79.
- Jacowitz, Karen 1995. Measures of Anchoring in Estimation Tasks, *Personality and Social Psychology Bulletin* 21: 1161-66.
- Johnson, Marcia and Carole Raye 2000. Cognitive and Brain Mechanisms of False Memories and Beliefs, in *Memory, brain and belief*, ed. Daniel Schacter, and Elaine Scarry, Cambridge: Harvard University Press: 35-86.
- Johnson, Marcia, Shahin Hashtroudi, and D. Stephen Lindsay 1993. Source Monitoring, *Psychological Bulletin* 114: 3-28.

- Kelley, Harold 1971. Attribution in Social Interaction, in *Attribution: Perceiving the Causes of Behavior*, ed. E. Jones, D. Kanouse, H. Kelley, R. Nisbett, S. Valins and B. Weiner, Morristown, NJ: General Learning Press.
- Kruglanski, Arie, and Donna Webster 1996. Motivated Closing of the Mind: “Seizing” and “freezing”, *Psychological Review* 103: 263-283.
- Kruglanski, Arie and Donna Webster 1991. Group Members’ Reactions to Opinion Deviates and Conformists at Varying Degrees of Proximity to Decision Deadline and of Environmental Noise, *Journal of Personality and Social Psychology* 61: 212-225.
- Kruglanski, Arie, Donna Webster and Adena Klem 1993. Motivated Resistance and Openness to Persuasion in the Presence or Absence of Prior Information, *Journal of Personality and Social Psychology* 65: 861-876.
- Kruglanski, Arie 1989. *Lay Epistemics and Human Knowledge*, New York: Plenum Press.
- Kruglanski, Arie and T. Freund 1983. The freezing and unfreezing of lay-inferences: Effects on impressional primacy, ethnic stereotyping, and numerical anchoring, *Journal of Experimental Social Psychology*, 19: 448-468.
- Kunda, Ziva 1990. The Case for Motivated Reasoning, *Psychological Bulletin* 108: 480-498.
- Lerner, Jennifer and Philip Tetlock 1999. Accounting for the Effects of Accountability, *Psychological Bulletin*, 125: 255-275.
- Lewis, David 1996. Elusive Knowledge, *Australasian Journal of Philosophy* 74: 549-567.
- Ludlow, Peter 2005. Contextualism and the New Linguistic Turn in Philosophy, in *Contextualism in Philosophy*, ed. Georg Preyer and Gerhard Peter, New York: Oxford University Press, 11-50.
- MacFarlane, John 2005. The Assessment Sensitivity of Knowledge Attributions, *Oxford Studies in Epistemology* I, 197-233.
- Mayselless, Ofra and Arie Kruglanski 1987. What Makes You So Sure? Effects of Epistemic Motivations on Judgmental Confidence, *Organizational Behavior and Human Decision Processes* 39:162-183.
- McAllister, Daniel, Terence Mitchell and Lee Roy Beach 1979. Contingency-Model for the Selection of Decision Strategies – Empirical Test of the Effects of Significance, Accountability, and Reversibility, *Organizational Behavior and Human Performance* 24: 228-244.
- Morton, Adam and Antti Karjalainen 2003. Contrastive Knowledge, *Philosophical Explorations* 6: 74-89.

- Nickerson, Raymond 1999. How We Know--and Sometimes Misjudge--What Others Know: Imputing One's Own Knowledge to Others, *Psychological Bulletin* 125: 737-759.
- Richard, Mark 2004. Contextualism and Relativism, *Philosophical Studies* 119: 215-242.
- Sanitioso, Rasyid and Ziva Kunda 1991. Ducking the Collection of Costly Evidence: Motivated Use of Statistical Heuristics, *Journal of Behavioral Decision Making* 4: 161-176.
- Schaffer, Jonathan 2005. Contrastive Knowledge, *Oxford Studies in Epistemology* 1: 235-271.
- Schaffer, Jonathan 2006. The Irrelevance of the Subject: Against Subject-Sensitive Invariantism, *Philosophical Studies* 127: 87-107.
- Sherman, Steven, Robert Cialdini, Donna Schwartzmann and Kim Reynolds 1985. Imagining can Heighten Or Lower the Perceived Likelihood of Contracting a Disease: The Mediating Effect of Ease of Imagery, *Personality and Social Psychology Bulletin* 11: 118-127.
- Stanley, Jason 2005. *Knowledge and Practical Interests*, New York: Oxford University Press.
- Tetlock, Philip and Richard Boettger 1989. Accountability: a Social Magnifier of the Dilution Effect, *Journal of Personality and Social Psychology* 57: 388-398.
- Tetlock, Philip, Jennifer Lerner and Richard Boettger 1996. The Dilution Effect: Judgmental Bias, Conversational Convention, or a Bit of Both?, *European Journal of Social Psychology* 26: 915-934.
- Tversky, Amos and Daniel Kahneman 1974. Judgment Under Uncertainty: Heuristics and Biases, *Science* 185: 1124-1131.
- Van Hiel, Alain and Ivan Mervielde 2003. The Need-for-Closure and the Spontaneous Use of Complex and Simple Cognitive Structures, *Journal of Social Psychology* 143: 559-568.
- Webster, Donna 1993. Motivated Augmentation and Reduction of the Overattribution Bias, *Journal of Personality and Social Psychology* 65: 261-271.
- Williamson, Timothy 2005. Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly* 55: 213-235.