

Epistemic anxiety and adaptive invariance

What enables us to recognize, intuitively, that a person knows a certain proposition? One of the challenges we face in seeking a general answer to this question comes from the great diversity of propositions a person might know. It's not obvious what common factors unite, say, knowing that there is no greatest prime number with knowing that there is currently a mild smell of soap in the room. At least each of those cases suggests a clear path: demonstrative proof and olfactory perception readily come to mind as good ways of making up one's mind. An additional challenge is posed by cases in which no single route to knowing seems best, cases in which one might rely on testimony of one kind or another, perception, memory, inference, or some combination of those and other resources. Many of the examples that have animated recent work in epistemology have had this multifarious character; there is no intuitively obvious single best way of knowing whether some particular bank branch will be open this Saturday, where one's car is currently parked, whether some flight will stop in Chicago, or whether a particular co-worker one didn't see in person was actually at work today.¹

At least when a question is settled by demonstrative proof, or by one's sense of smell, there seems to be a natural stopping point for inquiry; in these more complex cases the natural endpoint is not so clear. Is it enough that I remember having seen John's hat hanging on the rack outside his office? Should I talk to someone who might have met with him, perhaps several people, or ask to review the security tapes from the building? Should I collect evidence on just how hard it would be for someone to tamper with those security tapes? There could be an open-ended range of cognitive activities one might pursue here, and it is not obvious how deep into the range a person would need to travel before attaining knowledge of the proposition at issue. As many epistemologists have observed, just thinking about the question of where to stop can make knowledge seem harder to attain. Under heightened self-consciousness, even those cases that initially came across as simple can start to seem harder to settle decisively: should I double-check the proof, or worry about my over-active imagination, or about the acuity of my sense of smell? Wondering about finding the right moment to close off inquiry and claim knowledge, one may feel haunted by Sextus Empiricus's suggestion that it is *always* possible to continue the search.

¹ The cases alluded to are in (Cohen, 1999; DeRose, 1992, 2009; Vogel, 1990).

The puzzle that will concern us in what follows is not so much about the difficulty of finding the right point to close off inquiry as it is about the disparity between the difficulty of that question considered in the abstract and the ease of intuitively recognizing the presence or absence of knowledge in particular concrete cases. ‘To know’ is one of our ten most common verbs (Davies & Gardner, 2010), and in daily life we rarely agonize over its application. But if it is easy in practice to get an intuitive sense that someone does or does not know that the bank will be open tomorrow, this is not because we commonly anticipate just the same behavior from any given subject who is making up his mind on that point. For example, it is widely observed that we expect subjects who know that much is at stake to think harder and collect more evidence before taking a question to be settled. In practice, it seems that there is something about the demands of thinking about a particular concrete case that helps give us a natural sense of how far a subject must go to gain knowledge.

In what follows I argue that the practical demands of a concrete situation will naturally generate some roughly predictable level of ‘epistemic anxiety’², a force that normally determines how much evidence we are inclined to collect and how thoroughly we will weigh it before making up our minds. This force is no mystery: the next two sections of this paper review a body of empirical work showing epistemic anxiety to be a natural aspect of the regulation of our thinking, a factor that works to ensure that cognitive activity integrates with other types of activity in balancing expected costs and benefits. Allocating more cognitive effort to a question generally leads to greater expected accuracy of judgment, but we have only so much effort available to allocate to all our endeavors. Where the expected costs of inaccuracy are high, we typically feel more of the epistemic anxiety that drives increased cognitive effort. In general, automatic variations in epistemic anxiety make our thinking more ‘adaptive’, so that epistemic behavior like evidence-collecting is governed in ways that tend to complement rather than hinder our other pursuits.

When we think about particular cases, we recognize that for some subjects, time spent collecting additional evidence could most rationally be spent otherwise. We get a clearer sense of whether or not that behavior is to be expected if we grasp the subject’s practical situation and the value she would naturally see in gaining additional accuracy on that question. If knowledge ascription involves the attribution of some psychological attitude like outright belief, the task of knowledge ascription should be expedited by our having an intuitive sense of when a person would

² For earlier uses of this expression see (Hawthorne, 2004; Hawthorne, 2005; Hookway, 1998). Hookway credits the notion to Peirce.

typically form an outright belief, or take a question to be settled. So, especially when we are considering complex cases where a proposition p would normally be known through the integration of a variety of cognitive resources, it would be easier to gain an intuitive sense about what is needed for knowledge of that proposition if we are considering not the abstract question of what it would take for any S to know that p , but a particular story about some S who aims to know whether p , where the practical circumstances of this S are well understood.

Using complex cases which readily produce variations in expected epistemic behavior, contextualists and relativists have diagnosed these variations in our expectations as arising from shifting epistemic standards. In what follows I argue that our varied expectations for epistemic behavior might be better seen as arising from an invariant expectation that subjects will think adaptively; in fact, I'll aim to show that we naturally see anxious high-stakes and casual low-stakes subjects as needing to engage in different behavior in order to meet precisely the same standard.

Given the emphasis it places on practical circumstances, the position I aim to defend – a position I'll call adaptive invariantism – might look like a form of subject-sensitive or interest-relative invariantism (henceforth IRI). Advocates of IRI have maintained that a subject's practical interests themselves figure as a factor in determining whether he knows, alongside traditional factors like the truth of the proposition he believes, his degree of confidence in that proposition, the reliability of his belief-formation method, and so forth (Fantl & McGrath, 2009; Hawthorne, 2004; Stanley, 2005). Granting that a representation of a subject's practical interests can help us gain an intuitive sense of whether this subject knows, I will argue that it would be premature to conclude that interests are themselves a knowledge-determining factor. If we look at the relationship between the traditional factors of reliability and outright belief, we can resist the IRI suggestion that a subject's practical circumstances themselves make a difference to whether his true belief intuitively counts as knowledge.

Section 1 of the paper examines some empirical work from the literature on judgment and decision making in support of the thesis that ordinary cognitive behavior is approximately adaptive, in the sense of balancing expected costs in cognitive effort against expected gains in accuracy. Section 2 describes the role of epistemic anxiety in governing these variations in cognitive effort, and argues for a close relationship between the satisfaction of epistemic anxiety and the normal formation of evidence-based outright belief. Section 3 defends adaptive invariantism as a view of intuitive knowledge ascription.

1. Adaptive cognition

Some cognitive problems can be handled in a variety of ways; this observation has interested psychologists as much as epistemologists. A great variety of empirical work has been done on variations in cognitive effort, and a broad consensus has emerged that these variations are roughly adaptive: other things being equal, people will typically invest more energy in making judgments when there is greater anticipated reward in the accuracy of those judgments, or greater anticipated cost in being inaccurate. This increased effort is generally rewarded with increased accuracy (for reviews, see e.g. Broder & Newell, 2008; Payne, Bettman, & Johnson, 1993).

Some early studies made the tradeoff between accuracy and effort very conspicuous to experimental participants. Daniel McAllister and colleagues invited their MBA student subjects to consider a variety of hypothetical business decisions that were either low-stakes (affecting 10% of available corporate funds) or high-stakes (affecting 80% of available funds). The students were asked to select between various more and less complex decision strategies (for example, applying more and less complex formulae for securing loans), with the more complex strategies explicitly marked as both more accurate and more costly in terms of tying up managerial resources; students had a restricted pool of managerial resource units available to cover all the decisions for which decision strategies were to be chosen. Unsurprisingly, the more complex and accurate decision strategies were seen as appropriate for the higher-stakes decisions (McAllister, Mitchell, & Beach, 1979, Study 1). Similar results were found when participants were asked not only to evaluate the hypothetical appropriateness of various strategies, but also to choose strategies themselves in solving problems presented in various cases marked as high- and low-stakes (McAllister et al., 1979, Study 2). More subtle experiments conducted in the following decades found high-stakes participants spontaneously gravitating to strategies involving greater cognitive effort even without being specifically cued to recognize those strategies as more accurate.

High-stakes subjects spontaneously collect more evidence before making up their minds. In a classic demonstration of this tendency, participants had the task of identifying a series of digits (1-9) flashed individually on a screen by a tachistoscope in very short (4 ms) bursts (Mayseless & Kruglanski, 1987). This timing had been established in pretesting to support 50% correct identification of a digit on a single exposure. After each scarcely visible presentation, participants were asked to at least provisionally guess the digit, and then to rate their confidence on a scale from 0 (“not at all confident”) to 100 (“confident beyond a shadow of a doubt”). Participants were allowed to activate the tachistoscope to see the brief flash of each digit again as many times as they

felt they needed to identify it; when a decision was reached, the next digit would start being presented, and so on until 10 digits had been named. Half the participants were in a high-stakes condition, promised an extra hour of (apparently much-desired) experimental credit for getting at least 9 out of the 10 digits right; the other half were low-stakes, asked to identify the digits but not promised any performance rewards for accuracy. High-stakes subjects collected more evidence (pressing the button on average just over 15 times, where low-stakes subjects were satisfied with an average of 4.75 presentations); high-stakes subjects also reported lower confidence in their initial estimations (reporting an average confidence of 30.27 against the low-stakes 40.75 following the first flash), but – unsurprisingly, given their greater evidence – substantially higher average confidence by the moment of settling the question on each digit (91.46 vs. 68.42).

For the digit identification task, differences in cognitive effort between low- and high-stakes subjects could be measured quite simply, by counting the number of flashes each participant chose to observe. For more complex problems, high and low stakes subjects can show variations in cognitive effort not only in collecting different amounts of evidence, but also in finding more and less taxing ways to weigh and integrate the evidence at their disposal. Various metrics have been proposed to measure the costs of different ways of collecting, integrating and weighing evidence, notably the “Elementary Information Process” scoring system popularized in *The Adaptive Decision Maker* (Payne et al., 1993). The maximal EIP value is given to the (in principle, maximally accurate) weighted additive rule, which requires subjects to survey all relevant cues, assess the weight of each, sum the weighted values of each cue for each alternative, compare all alternatives and choose the one with the highest value. Simpler heuristic strategies would permit a decision maker to “tally” or weight all cues evenly, or to consider fewer cues – perhaps even deciding on the basis of the first cue that discriminates between alternatives. Payne, Bettman and Johnson themselves were careful to underscore the limitations of the EIP metric, noting that individuals could differ somewhat on which aspects of weighing and integrating information they found most taxing, and that even simple functions such as adding could be more or less taxing depending on task conditions. With these cautionary restrictions in mind, they nevertheless collected evidence from a very wide range of empirical studies to show a clear association between high stakes, increased cognitive effort and accuracy.

Payne, Bettman and Johnson’s findings on preferential choice (or judgment about what to do) fit into a larger pattern of results on human judgment in general. For example, similar findings have emerged within the dual process theory (DPT) research program. The DPT program has

investigated the division between automatic, heuristic cognition and more controlled, analytic cognition (for reviews, see Evans, 2007; Frankish & Evans, 2009; Sloman, 1996). Heuristic cognition is involuntary and effortless, where analytic cognition requires attention and effort; indeed, one way of organizing the various kinds of heuristics into a single framework is to classify them in terms of the type of reduction in cognitive effort that each provides (Shah & Oppenheimer, 2008). Stakes have been identified as one factor that will curb heuristic cognition by increasing cognitive effort (e.g. Epley & Gilovich, 2005).³

There are several different ways of conceptualizing the variations in cognitive effort that are agreed to accompany shifting stakes. On the ‘adaptive toolbox’ approach, various features of the decision environment trigger the use of a rougher or finer cognitive strategy; notably, we are inclined to adopt heuristic strategies that will involve considering less evidence when motivation is low and the cost of collecting evidence is high (Gigerenzer & Todd, 1999). Advocates of the adaptive toolbox approach contend that we have a broad repertoire of these strategies, some more involved and taxing than others. Meanwhile, on the rival ‘evidence accrual’ approach, what gets determined by the decision context is not a particular cognitive strategy, but rather an evidence threshold; evidence is always ‘sampled’ sequentially until one alternative has sufficient support to pass this preset threshold (Lee & Cummins, 2004; Newell, 2005). Differences between higher and lower evidence thresholds can be observed when, for example, the leading piece of evidence tells in favor of option A, but the entire body of evidence ends up telling in favor of option B. In such settings, subjects who are more motivated – for example, given larger financial rewards for accuracy – search for and integrate more evidence and are more likely to select option B (Newell & Lee, 2010).

There is no consensus about whether adaptivity is better understood in terms of the adaptive toolbox or the evidence accrual model.⁴ It also remains controversial exactly *how* adaptive we are, how well our epistemic behavior is tailored to the pursuit of our goals. While there is a general pattern of seeking more information for higher-stakes decisions (and less information

³ More analytic cognition also emerges under high accountability conditions, in which one anticipates having to justify one’s judgments to others (Lerner & Tetlock, 1999; Simonson & Nye, 1992). McAllister and colleagues also found both the expectation and the expenditure of greater cognitive effort for high-accountability decisions, and for decisions that were marked as irreversible. (McAllister et al., 1979). Making a decision irreversible is arguably a straightforward manipulation of stakes; the relationship between accountability and stakes is more complex, and will be discussed in Section 3.

⁴ In their 2008 review of the literature, Arndt Bröder and Ben Newell concluded that ‘At the moment, data do not allow for a clear decision between the model classes because apparent strategy switches can be reinterpreted as criterion shifts or vice versa.’ (Broder & Newell, 2008, 209) They do however outline some of the issues in strategy selection learning, performance, and theoretical simplicity whose investigation should favor one type of model over the other.

when information-gathering is expensive), there is considerable individual variation in evidence collection on the same task (Newell & Shanks, 2003). There are also some apparent deviations from perfect adaptivity that are more systematic, although it is not yet clear whether these are more than apparent. In the artificial task environments in which the price of evidence can be manipulated and the level of evidence collection can be measured with precision, subjects display some systematic tendency to under-purchase information when it is cheap relative to monetary payoff, and over-purchase when it is expensive (Connolly & Serre, 1984; Newell, Weston, & Shanks, 2003). However, it is possible to interpret these apparent failures of calibration as manifesting the rational responses to costs and benefits other than the explicit information charges and accuracy payoffs within the experiment. There are costs to collecting evidence beyond the monetary charge for each piece of information accrued; it is tiresome and time-consuming to incorporate larger ranges of evidence in judgment (Holland & Kleinmuntz, 1994). Meanwhile, there are also benefits to information beyond its contribution to the likelihood of winning monetary rewards for accuracy. Subjects in environments with high charges for evidence relative to payoff generally do respond to these charges by purchasing less evidence, but their information purchase is not curbed as far as it should be to maximize financial return. However, as one experimental team noted, it is possible that these subjects might be placing some value on gaining knowledge for its own sake (Newell, Rakow, Weston, & Shanks, 2004, 130). Given that they were working with only a few pieces of relatively costly evidence, they would not have felt the same information burden as their counterparts overwhelmed with cheap evidence.

If it is difficult to assess deviations from perfect adaptivity in artificially controlled settings, it is harder still to gauge just how adaptive our thinking is in facing real-world problems. Fortunately, for our present purposes it may not matter precisely how adaptive we are in changing our cognitive strategies or evidence thresholds: the well-established global pattern of approximate adaptivity is enough to explain our responses to stark contrasts like those of DeRose's Bank Case and Cohen's Airport Case. We can appreciate that the subject who faces a very serious financial penalty if the bank is closed tomorrow would be more likely to double-check the hours than the subject who faces at most the minor inconvenience of a wasted trip.⁵

In considering our intuitions about these cases, one might be concerned about the switch from the first- to the third-person point of view. It is one thing to show that we naturally respond

⁵ The question of the exact quality of our calibration will be re-opened in the last section of the paper.

to our own increased stakes by gathering more evidence and processing it more thoroughly; it could seem to be quite another matter whether our attributions of knowledge and belief reflect the expectation that others will think in even a roughly adaptive manner. To address this worry, one might begin by observing that the adaptive cognition literature has itself covered both types of assessment: the early studies done by McAllister and colleagues, for example, generated similar results when they asked subjects what cognitive strategies they would anticipate in another person with high or low stakes, and when they manipulated their participants' own stakes (McAllister et al., 1979).

The reason for this uniformity arguably lies deep in the mechanics of intuitive propositional attitude attribution. The dominant models of propositional attitude attribution both posit a very significant overlap between the conditions governing our own mental states and the conditions we see as relevant to the mental states of others. On the simulation model, one imaginatively puts oneself in the position of another, pretending to have his beliefs and desires, then makes whatever judgment would be the natural outcome of those pretended initial states, and finally projects the resultant state back onto the third party (Goldman, 2006). In the “theory-theory” model, adult propositional attitude ascription involves the application of naïve theories, where these theories come to incorporate an extensive track record of experience about the relationship between mental states and behavior, both for self and others (Gopnik & Meltzoff, 1997; Perner, 1991). Neither of these models predicts a perfect congruity between first- and third-person mental state ascriptions – in particular, both models end up predicting, in different ways, that we will have difficulty reading the mental states of those who know less about their predicaments than we do, a finding that will be reconsidered in connection with the discussion of Ignorant High Stakes cases in Section 3. However, both models would agree that if large changes in stakes generally have a significant impact on how we think and behave, then adult intuitive assessments of others should reflect this variation.⁶

⁶ One recent paper in experimental philosophy purports to show that stakes do not influence knowledge ascriptions (Schaffer & Knobe, forthcoming). However, the experiments in this paper were not set up to ensure that participants would construe the high- and low-stakes counterparts as having equal evidence; other things being equal, we anticipate higher-stakes subjects would have more evidence, given their level of motivation, and this natural expectation is enough to explain the null result obtained. When care is taken to ensure that attributed evidence is held fixed across the contrast, higher stakes do seem to increase stringency in knowledge ascription (Pinillos, submitted).

2. Epistemic anxiety and outright belief

Sometimes, belief formation is automatic and effortless; at other times – and notably, in high-stakes circumstances – it can be harder to make up one’s mind. To remain neutral between the adaptive toolbox and the evidence accrual approach, it will be useful to introduce a generic expression for the heightened need for greater evidence and more thorough processing that is characteristic of high-stakes situations. In what follows I will use ‘epistemic anxiety’ as the generic label for the inclination or desire for increased cognitive activity, and I will try to remain neutral about whether we should think of epistemic anxiety as setting a higher evidence threshold or selecting a more demanding strategy from the toolbox. Whichever way we think of it, epistemic anxiety can either be satisfied by amassing and weighing evidence, or it can be overridden by other desires, in ways to be discussed in what follows.⁷ Where a question elicits epistemic anxiety, we can achieve outright belief on that question either by satisfying our epistemic anxiety (culminating in normal evidence-based belief), or by overshadowing that anxiety (as happens in cases of wishful thinking, and under conditions of haste and distraction). This section aims to explain the relationship between epistemic anxiety, cognitive effort and better or worse kinds of outright belief formation.

Any discussion of cognitive effort in the formation of outright belief may run the risk of suggesting an implausible voluntarism. As a start towards allaying those concerns, it should be emphasized that one does not typically choose one’s level of epistemic anxiety. The expression ‘strategy selection’ is often used to refer to the adoption of a particular heuristic or analytic way of solving a problem, but the ‘selecting’ here is not typically achieved by personal-level conscious reflection on the merits of the various alternative ways of thinking. Indeed, as Jörg Rieskamp and Philipp Otto among others have noted, if strategy selection always had to be a matter of conscious decision, one could well be stuck with ‘a recursive homunculi problem of deciding how to decide’ (Rieskamp & Otto, 2006, 207). Advocates of the adaptive toolbox approach therefore see strategy selection as handled in the typical case by unconscious processing of relevant features of the decision environment. Similarly, within the evidence accrual model, a higher or lower ‘preset’ evidence threshold would typically be computed automatically on a given task; the task of setting the threshold is not itself typically a matter of conscious decision, again on pain of regress (Newell,

⁷ Epistemic anxiety about some proposition p is not always satisfied or overridden; it can also fade away, say when p for some reason ceases to matter. Because such cases involve no judgment, they are not particularly central to recent debates about the difference between belief and knowledge ascription. There are also cases in which one reaches suspension of judgment, either on the basis of reflection on the evidence, or on some other basis. These cases merit close attention, but will be set aside for present purposes.

2005). Either way of modeling the problem can take the setting of epistemic anxiety to be itself something that typically incorporates some combination of innate dispositions, current stimuli (including metacognitive feedback), and learning from past decisions (Alter, Oppenheimer, Epley, & Eyre, 2007; Gigerenzer & Todd, 1999; Newell, 2005; Rieskamp & Otto, 2006). For present purposes our focus is on stakes as a source of increased epistemic anxiety, but it is worth mentioning that epistemic anxiety can also be increased in other ways. It can matter how a question is framed, for example: questions posed in familiar terms can trigger automatic ‘overlearned’ responses, where unfamiliar questions demand controlled, step-by-step reasoning.⁸

There are benefits to having automatic processes set varying thresholds for the conversion of evidential support into outright or binary belief. Some benefits can be appreciated if we think of outright belief as the solution to a problem of bounded rationality. One connection between outright belief and bounded rationality is often noted: perhaps with unbounded time and cognitive powers, we could always maintain exquisitely subtle attitudes towards propositions, reflecting our precise level of evidential support for each; given the limits on our resources, we may be better served by a rougher capacity to say yes or no, a capacity that provides simpler binary inputs to future inferences and decisions (Foley, 1993, 2009; Harman, 1986). Binary states of outright belief may be needed as inputs to judgments because of limits on our cognitive resources, but there is a further problem in bounded rationality having to do with the initial production of these binary states of outright belief. It is adaptive for us to invest more effort in the formation of outright belief in high-stakes circumstances, but we cannot reflect indefinitely on the level of energy to invest in any given judgment. By having automatic processes set varying levels of epistemic anxiety, we are placed in the desirable middle ground between a fixed-effort system of judgment in which the same level of epistemic activity would be applied to all judgments and a hyper-deliberate system in which we would have to reason explicitly about which stopping rule to impose on evidence collection on each occasion.

Once a subject’s level of epistemic anxiety is set for a certain question, his desire to expend cognitive effort on that question will compete with all the other desires he is currently

⁸ To use an example of Keith Frankish’s, “What is 9 times 11?” is answered automatically, and “What is 9 divided by 11?” requires effortful sequential thought (Frankish, 2009). Controlled cognition is also called into play to negate hypothetical possibilities. In normal conditions, we would give an automatic, heuristically-generated answer to the question, “Do you know who is the current President of the USA?”, but we would have to engage in controlled cognition to answer the question, “Do you know that Obama has not had a heart attack in the last five minutes?” For further discussion of the epistemological significance of this contrast, see (Nagel, forthcoming).

experiencing. Where epistemic anxiety has sufficient relative strength, or perhaps where it is consciously evaluated as worth pursuing, the subject will act to satisfy it, for example, by collecting evidence. Collecting and weighing evidence is something that takes effort, but again, this is not the sort of effort that would need to have a problematically voluntarist character: in the normal case controlled evidence-acquisition is not driven by a desire to believe that p (or by the desire to believe that not- p). Epistemic anxiety is itself equally well satisfied whichever alternative ends up passing the threshold or being selected by the strategy.

Once worries about problematic sorts of voluntarism have been addressed, the type of control afforded by epistemic anxiety can be better appreciated. Increased cognitive effort is not produced immediately by any high-stakes situation; because the consciously accessible factor of epistemic anxiety mediates the relationship between triggering factors in the decision environment and increases in cognitive effort, it becomes possible for us to weigh the importance of satisfying that anxiety against the importance of satisfying other pressing desires in our conscious deliberations.

This generic notion of epistemic anxiety would not be the only consciously accessible signal that has a part in regulating cognitive conduct; some comparison with other such signals may be illuminating. Some of the better-known regulatory signals are metacognitive feelings, feelings that deliver information about the operation of our own cognitive processes. Most notably, metacognitive feelings deliver information about the fluency or ease of operation of various underlying cognitive processes. When we retrieve a fact from memory, for example, we become aware not only of what is retrieved, but also of our ease of retrieval; in fact, our confidence in what is retrieved is often inferred from our fluency in retrieving it (Kelley, 1993; Schwarz, 1998). It is thought that feelings of fluency form the basis of subjective confidence across a great range of different types of cognitive operation, including perception, spatial reasoning, linguistic processing, memory and deduction (for a comprehensive review, see Alter & Oppenheimer, 2009).⁹

In the particular case of memory retrieval, one noteworthy metacognitive signal is the FOK or *feeling of knowing*, a signal we experience in its most vivid form as the tip-of-the-tongue feeling, occurring when an answer seems close but retrieval is blocked. The FOK is not a meaningless signal: for subjects unable to generate an answer to a fill-in-the-blank test on a given question,

⁹ There is also evidence that where experienced fluency falls short of what was antecedently expected, epistemic anxiety is increased (Alter et al., 2007).

strength of FOK correlates with the likelihood of answering that question correctly on a subsequent multiple-choice test (Hart, 1965). The FOK is not generated by any direct access to the hidden target, however; rather, it is produced by monitoring the fluency of recall, and in particular by the volume of cue-associated information that is swiftly brought to mind (Koriat, 1993, 1995). Not all failures of recall generate an FOK: where we do experience a tip-of-the-tongue feeling, it typically motivates us to invest more cognitive effort in searching memory for an answer. Given that FOK is generated by the accessibility of information associated with the target, and the possession of this associated information is often correlated with possession of the desired answer, this signal to continue searching memory is largely adaptive (Koriat, 2000; Koriat & Goldsmith, 1996).

Asher Koriat and Ravit Levy-Sadot argue that certain practical advantages accrue to having cognitive effort guided by metacognitive signals that are accessible to the subject: 'when unconscious influences give rise to a subjective, conscious feeling, that feeling can be integrated with the goals and other beliefs of the person, and modulated by them in guiding judgment and action. The result is that judgments and actions are more intelligent and adaptive than when they represent the product of unopposed, unconscious and automatic influences.' (Koriat & Levy-Sadot, 2000, 199) On this account, the FOK is not simply an epiphenomenal conscious accompaniment of blocked recall, but a signal whose importance can be weighed against our other priorities. Given the limits on available time and effort, we can decide whether the cost of searching our memory in response to the FOK is a worthwhile expenditure.

The function of epistemic anxiety then runs roughly parallel to the function outlined for the FOK: epistemic anxiety is not simply an epiphenomenal accompaniment of high-stakes situations. High epistemic anxiety is a signal that the task of settling the question *whether p* on the basis of evidence will take significant effort. The conscious availability of epistemic anxiety places the problem of whether to undertake this task among the other problems falling under the scope of practical reason.

The larger goal here is to show how variations in perceived epistemic anxiety can explain certain patterns in our attribution of knowledge, in particular, patterns involving changing stakes. If we need to see subjects as having made up their minds on the basis of evidence in order to see them as knowing, then changes in stakes could make a difference to whether a subject seems to know simply by making a difference to whether they seem to have made up their minds. It is beyond the scope of this paper to give a full account of the relationship between epistemic anxiety and making

up one's mind (or settling the question *whether p*, or forming an outright belief—I use these expressions interchangeably), and this is at least in part because the nature of outright belief is itself a controversial and difficult topic. But for our purposes here, it would be enough to establish that someone with high perceived epistemic anxiety would need to be seen as having more evidence than his low-anxiety counterpart in order to be seen as having a normal evidence-based outright belief. Here a critic might agree that the anxious person will keep collecting evidence longer than his casual counterpart, while still thinking that this anxious person would nonetheless be seen as already having an outright belief as soon as he has the moderate level of evidence possessed by his counterpart. This worry is hard to answer without taking on the task of defending some particular view of what it is to have an outright belief. As an interim measure, it is possible to look at several popular—and quite different—accounts of what it is to have an outright belief, and verify that on each of them, high perceived epistemic anxiety would indeed be correlated with perceived difficulty in attaining outright belief.

Functionalist conceptions of belief serve as one possible place to start, not least because functionalism has worked as a motivation for one prominent view of the relationship between stakes and outright belief, a view advanced by Brian Weatherson. Taking inspiration from “the functionalist idea that to believe that *p* is to treat *p* as true for the purposes of practical reasoning” Weatherson argues that stakes make a difference in the formation of outright belief because the same degree of confidence in a proposition might in some practical environments but not others rationally license a subject's treating that proposition as true (Weatherson, 2005, 421). It is consistent with that position to note that if practical reasoning includes in its scope all controlled action, even epistemic action such as evidence-seeking, then the epistemically anxious person who still wants to allocate effort to collecting evidence on the question of *whether p* is not simply treating *p* as true for the purposes of practical reasoning (cf. Ganson, 2008).

The Weatherson view of the relation between degree of confidence and outright belief fits neatly with much of the empirical work described in the last section. Certain psychological accounts of epistemic anxiety even use terminology similar to Weatherson's: Daniel Hausmann and Damian Lage, for example, describe the variable evidence threshold as marking a “desired confidence level”, which rises and falls in step with stakes, among other factors (Hausmann & Lage, 2008). Alice Eagly and Shelly Chaiken describe transitions between heuristic and systematic ways of settling questions as motivated by discrepancies between the subject's actual and desired confidence levels,

where higher-stakes problems set higher desired confidence levels (Eagly & Chaiken, 1993). In Ofra Mayselless and Arie Kruglanski's work, subjects called off their search for evidence and reached a final decision on each digit at noticeably different average levels of reported confidence when they were in low-stakes as opposed to high-stakes conditions – 68.42 vs. 91.46 on the 0-100 scale (Mayselless & Kruglanski, 1987).

Weatherson does not assign any independent psychological reality to the notion of outright belief; what counts as an outright belief simply is a certain level of confidence, where only the confidence itself is psychologically real. Various considerations might drive one towards a more robust view of outright belief. To begin, one might think that outright belief has psychologically significant phenomenal qualities. Where high degrees of confidence are compatible with a continuing feeling of uncertainty, outright belief yields the subjective sense of a solid result. Psychologists who have studied these phenomenal qualities have characterized them in a way that would support the notion that outright belief formation marks the end of epistemic anxiety: we switch from an anxious phase of searching for evidence and weighing it to a satisfied or relieved phase of possessing a determinate representation (Gollwitzer & Bayer, 1999; Kruglanski & Webster, 1996).

Various models of judgment posit a psychologically significant and effortful transition between a processing stage and a result stage, where there are substantive differences between these stages, and motivation is needed to produce the transition between them. For example, in Evans's heuristic-analytic theory, reasoning in high-stakes circumstances requires an effortful examination of various mental models, where in low-stakes circumstances we could accept the most plausible relevant model automatically. On this account, unsatisfied epistemic anxiety would ordinarily keep us engaged in the project of constructing and evaluating various mental models until we find one that satisfies or passes our threshold. Whether we can make judgments and inferences that take p as input depends not on our level of confidence in p but on whether we have yet accepted a model in which p holds (Evans, 2006). Like the functionalist account, a more robust account like Evans's can accept that the satisfaction of epistemic anxiety culminates in the normal formation of evidence-based belief. Different accounts of outright belief will have different ways of describing its relationship to epistemic anxiety, but it is noteworthy that one of the points on which varied accounts of outright belief seem to agree is that high epistemic anxiety would bar the normal formation of outright belief on moderate evidence.

So far the focus has been the normal formation of evidence-based belief, but not all belief formation is normal. When unopposed by other forces, epistemic anxiety will spur cognitive effort, plausibly terminating in evidence-based outright belief. But other forces can intervene, and outright belief can be formed by processes other than the normal evidence-based ones. Sometimes—for example, in certain cases of wishful thinking—it is pragmatically rational to curtail epistemic effort, and in these cases a desire other than the desire for evidence can govern our epistemic behavior and lead us to settle a question one way rather than another.¹⁰ This might be thought of as the dark side of epistemic anxiety: because it places epistemic behavior under the scope of practical reason, epistemic anxiety creates an opening for our pragmatic interests to override our epistemic interests in those cases where these interests are opposed. Meanwhile, even when our desire for evidence is not competing with any desire to conclude a question one way rather than the other, conditions such as distraction can quite generally impair our capacity to allocate effort in accordance with our desires. It might be rational for me to devote more controlled cognition to some high-stakes question, but under cognitive load I may have insufficient resources to meet this rational demand, and make up my mind in some less rational manner.

On the picture sketched so far, ordinary subjects automatically experience levels of epistemic anxiety roughly reflecting their stakes (and certain other factors, some of which will be discussed in the next section). In normal circumstances, subjects will respond to increased anxiety by thinking harder. However, factors that stop or distract us from satisfying epistemic anxiety can enable us to make up our minds prematurely. Now, if we ordinarily make sense of others by ascribing beliefs and desires to them, our intuitive assessments of the epistemic positions of others should take these natural patterns into account. For example, if someone is in a high-stakes situation and declines to pursue readily available evidence on a question that should be provoking high epistemic anxiety, it would be natural for us to attribute to him some desire or condition overshadowing his natural desire for increased cognitive effort. If we see this condition as the basis of his belief, then

¹⁰ In some extended sense of ‘evidence-based’, perhaps even wishful thinking should count as evidence-based belief formation. In a review of empirical work on motivated reasoning, Ziva Kunda concludes that in wishful thinking we do attend to evidence, but in a very asymmetrical way, devoting much more attention to evidence supportive of our desired conclusion (Kunda, 1990). If Kunda is right, wishful thinking might simply influence the direction of controlled cognition rather than the amount of it: cognitive effort could be high where anxiety is high, but distorted in its character by the desire to reach one conclusion rather than the other. For present purposes I will use the expression ‘evidence-based’ to pick out belief formation that is not distorted in this manner, but with some rewording the general point of this section could equally well be made within Kunda’s framework: we would naturally posit an accuracy-compromising desire to explain the distorted cognitive behavior of wishful thinkers.

his judgment may naturally seem less reliable than the judgment of his low-stakes counterpart.

The reliability of a subject's belief formation is widely agreed to make a difference in intuitive knowledge ascription, for traditional and non-traditional epistemologists alike. However, we face well-known difficulties both in individuating the process whose reliability is at issue, and in setting any particular threshold of reliability as sufficient for knowledge. The next section explores a hypothesis about the reliability factor in knowledge: perhaps what we look for in the reliability condition takes into account the level of epistemic anxiety we automatically feel on behalf of the subject we are evaluating.¹¹ To come across as knowing, subjects who are settling a given question – say, whether the bank will be open – need to collect a level of evidence that would satisfy the epistemic anxiety we naturally ascribe to them. Failure to do so would make it natural for us to see the subject as suffering from an accuracy-compromising condition such as haste, distraction or wishful thinking. Subjects come across as passing the reliability condition if they are seen as forming evidence-based beliefs, where more or less evidence may need to be collected in order to seem to satisfy that condition.

On this view the anxious high-stakes subject may need to visit the branch to double-check the banking hours in order to seem as accurate as the casual low-stakes subject who just remembers a recent Saturday trip to the bank; if the high-stakes subject doesn't work harder than his counterpart, he will come across as having a problematic basis for his belief. Although we expect different behavior of these two subjects, we are not holding them to different standards: as far as reliability is concerned, both subjects are benchmarked against the performance that would be anticipated if nothing opposed the level of epistemic anxiety that ascribers automatically feel on their behalf. What matters is not the number of pieces of evidence subjects are collecting to settle a question, but the type of belief formation they seem to display. To contrast the relevant types here – evidence-based reasoning on the one hand; wishful, distracted or hasty thinking on the other – we do not need to make use of resources that lie beyond the scope of traditional (invariantist, intellectualist) epistemology. Invariant prohibitions on biased and careless thinking can be applied across the board in sorting knowledge from mere belief.¹²

¹¹ Caveat: in cases where we have more information about the subject's predicament than he does (Ignorant High Stakes, Gettier cases), we may feel more epistemic anxiety than he would. Because we have some interesting limitations in evaluating the mental states of more naïve agents, these cases introduce new complexities, and will be tackled only after the more basic cases have been discussed.

¹² It is an open question whether subjects who are seen as biased and careless are judged not to know because they are perceived as unreliable, or whether the recognition of epistemic "vices" such as bias directly

3. Adaptive invariantism

Much recent non-traditional epistemology has focused on pairs of stories involving subjects who are supposed to come across as equally well-positioned with respect to all the factors that matter to traditional epistemology. Only some non-traditional factor is supposed to distinguish these matched counterparts. We are then invited to either make or consider natural epistemic evaluations of the twinned subjects, and where we find a contrast between these natural evaluations – say, that the low-stakes subject knows and the high-stakes subject does not know – we are urged to take this contrast as a sign that non-traditional factors matter to intuitive knowledge ascription. The general strategy of this section will be to argue that the adaptive invariantist can grant that our intuitive responses to these cases are different, but point to perceived differences in traditional factors to explain why. These differences may have escaped notice because they are automatically computed rather than explicitly stipulated in the scenario: when it is time to analyze our responses to the scenarios it is easy for us to recognize that the paired subjects have the same pieces of information in their initial evidence, and perhaps harder to notice that we have automatically attributed different levels of epistemic anxiety to them in response to the information about their stakes, with consequent changes in our expectations for the relationship between that evidence and their state of belief.

In the most heavily cited examples, perceived stakes are the non-traditional factor distinguishing each subject from his counterpart.¹³ If the argument of the last two sections is correct, this difference in perceived stakes would produce a difference in the level of epistemic anxiety one would automatically ascribe to the paired subjects. Because one's level of epistemic anxiety can make a difference to the way in which one bases one's belief, subjects who have the same evidence but different perceived levels of epistemic anxiety are not necessarily represented as forming their beliefs in the same way. If the traditional epistemologist is allowed to help herself to the premise that normal human cognition is expected to be adaptive, she may be able to explain

blocks knowledge ascription without the mediation of any judgment of reliability. Either possibility could be congenial to the traditional intellectualist invariantist. Thanks to Jonathan Weisberg here.

¹³ In some cases – Vogel's car theft scenario, Cohen's furniture store scenario – the crucial difference is not stake-related. I think these cases are also set up to induce changes in the level of epistemic anxiety experienced by the reader as she shifts between evaluating the paired subjects, but I will not argue that point here. Some relevant argument is presented in (Nagel, 2010, forthcoming).

the key cases motivating non-traditional epistemology. Some detailed discussion of these cases is in order, starting with the easier first-person cases.

The paired subjects in Keith DeRose's first-person Bank Cases – for ease of exposition we'll call them LOW and HIGH – are making judgments about a bank being open on Saturday. Each of them is outside the bank on a Friday, hoping to postpone a bank visit for tomorrow, and responding to a query from his wife about the possibility that the bank will be closed the next day. LOW, who has little at stake, seems right to tell his wife "I know it will be open" on the strength of his memory of a recent visit – indeed he cites his memory of this recent trip as his evidence in responding to his wife, who has mentioned that many banks are closed on Saturdays. Meanwhile HIGH also mentions a memory of a recent Saturday trip, but his wife then raises a more specific worry that banks sometimes change their hours. Reminded that he will be in serious trouble if he makes a mistake on the issue, HIGH seems equally right to say "I don't know. I'd better go in and make sure."

On an adaptive invariantist reading of these cases, HIGH is naturally seen as more motivated to research the question of whether the bank will be open. When he mentions his decision to search for more evidence, the mental state that would most easily make sense of this behavior is a state of being in the process of making up his mind on this issue. (And to make sense of LOW's choice to deflect his wife's suggestion that they stop at the bank now, we would naturally ascribe to him an outright belief that the bank is open tomorrow.) But if the question is not settled for HIGH, then there is a very traditional reason why he lacks knowledge: he lacks outright belief. A memory of a recent visit might be enough to secure outright evidence-based belief in a low-anxiety setting (indeed, by mentioning his reason together with his self-ascription of knowledge, LOW draws attention to the evidence-based character of his belief). In a high-anxiety setting, particularly one in which additional evidence is readily available, and we have been reminded of hypothetical situations in which memories of recent visits would not be decisive, we expect a subject to work harder. The fact that HIGH has just announced his intention to work harder fits these natural expectations, further supporting our sense that he is still making up his mind.

DeRose himself takes the psychological attitude of the paired subjects to be held fixed by a stipulation that HIGH remains as confident as he was before that the bank would be open. It is arguably somewhat difficult to register this stipulation, given the other content in the scenario: in announcing a decision to 'go in and make sure', HIGH certainly seems to be displaying lowered

confidence, in some sense of that term.¹⁴ Perhaps with effort, however, we can find some psychologically realistic ways of reading this stipulation. Most directly, we could read it as a stipulation that HIGH and LOW would have the same credence or subjective confidence in the proposition that the bank will be open. But this reading would not necessarily settle the question of whether HIGH has an outright belief, if in fact different levels of subjective confidence are required for outright belief in high- and low-stakes circumstances, as Weatherson has suggested. Independently, if HIGH is seen to have just the same subjective confidence as his counterpart, this could be a sign of old-fashioned trouble in HIGH's epistemic predicament: ordinarily, as we actively consider a broader range of hypotheses consistent with our evidence, confidence in any one of these hypotheses should fall (Kelley et al., 1972). Failure to experience diminished confidence about the bank being open in the face of new worries could naturally be taken to signal some force offsetting that natural drop, such as a pre-existing desire to believe that the bank will be open, perhaps motivated by immediate reluctance to stand in the long Friday line. With such a force in play, it is no longer clear that our anxious HIGH is twinned with causal LOW on all the factors that matter to traditional epistemology. Similar concerns arise if we understand the stipulation of continuing confidence as a stipulation that HIGH maintains outright belief throughout the scenario; and in addition, some new concerns could be generated this way.

If HIGH maintains full outright belief that the bank will be open even as his wife raises the concern about a change in hours, then it is not an open question for him whether the bank will be open, even as he announces his decision to "go in and make sure". If HIGH's mind is already made up, he is only pretending to feel the need to go in and check the hours.¹⁵ If we sympathize with the wife's concerns (given what is at stake, HIGH should expend the extra effort before making up his mind), we should be troubled by HIGH's refusal to take these concerns seriously. Again, if we need to see a subject as having some resistance to evidence-collection in order to explain his mental state, then there would ordinarily be something wrong with the basing of that state. On the other

¹⁴ Patrick Rysiew notes that the stipulation of equal confidence might not automatically register in the manner DeRose hopes (Rysiew, forthcoming); Kent Bach also worries that it is hard to reconcile with the other content of the scenario (Bach, 2005). Keith DeRose himself observes this problem, but argues that in the sense of 'confidence' most relevant for knowledge ascription there is a way of taking the stipulation seriously (DeRose, 2009, 191). He then outlines this sense of confidence as something like subjective probability or degree of belief; however, he has not yet endorsed any particular account of the relationship between degree of belief and outright belief. Without guidance on the relationship between the degreed notion and the binary one, it is not clear whether DeRose's notion is in fact what we care about when ascribing knowledge.

¹⁵ Or perhaps his real concern is to reassure himself that he knows; the pursuit of iterations of knowledge will be discussed shortly.

hand, if we do not sympathize with the wife's concerns, we could see HIGH as perfectly reasonable in maintaining outright belief throughout the conversation, but then we could also see him as retaining knowledge throughout. On this way of reading the case, HIGH would not be speaking sincerely in denying that he knows the bank will be open, even if it would be pragmatically rational for him to deny that he knows in order to placate his wife. We could find his denial of knowledge acceptable on this reading, but most easily by registering it as pragmatically rational rather than literally correct. In summary, however we understand the confidence stipulation – and it is already a worry that it can be understood in various ways – there is reason to worry that the first-person Bank case does not provide a clear or uncontroversial example of a shift in epistemic standards in response to non-traditional factors.

Before tackling the stronger third-person cases for shifting standards, it is worth addressing an objection to the first and I think most natural adaptive invariantist strategy, according to which HIGH lacks knowledge because he lacks outright belief. John Hawthorne notes that while it is tempting to think that salience of error possibilities might 'destroy knowledge that p by destroying belief that p ', there is nonetheless 'a perfectly reasonable sense of "belief" in which one believes (...) even when the possibility of error is salient in the relevant sense.' (Hawthorne, 2004, 169) The observation is hard to deny: even as HIGH gets out of the car to check the hours, there is indeed a sense in which he could naturally be described as believing that the bank will be open the next day. He considers it likely; he would be surprised to find out that the hours had in fact changed. If he were forced to choose, one way or the other, he would say that the bank would be open rather than closed; the bank's being open is what he would automatically represent as 'the single most plausible state of affairs', in Evans's terminology. Because his high credence in the proposition that the bank will be open would ordinarily support the outright judgment that the bank will be open, he can fairly be described as having a general tendency to judge that the bank will be open, and for many purposes having a general tendency to judge that p would count as believing that p .

However, if we attend more closely to HIGH's predicament just as he is double-checking the hours, it seems most accurate to say this tendency of his has been suspended: if the categorical basis of his dispositional belief is his representation of the bank's being open as the most plausible outcome, he is now in a situation in which his judgments are no longer automatically produced by what he represents as the most plausible outcome. HIGH's increased cognitive anxiety generates attention to additional hypotheses about other possible outcomes, such as the explicitly stated

hypothesis that the bank has changed its hours, and until he checks, he lacks evidence to rule these out. If HIGH is really taking it to be an open question, and not just going through the motions of checking, then he does not at that moment fully accept the proposition that the bank will be open, so there is a significant and saliently action-guiding sense in which he does not believe that it will be open. To be sure, if he were forced to choose, he would say the bank will be open, but it is important that he is not now being forced to choose; the absence of immediate pressure enables him to suspend his belief and collect more evidence. So, thinking about HIGH's general outlook we might describe him as a believer; thinking more closely about his immediate predicament, we can see him as someone whose state of belief is in suspension.¹⁶ In being asked to ascribe mental states that will make sense of that immediate behavior, we are naturally focused on his current predicament of suspension rather than his general tendency towards outright belief.

Some further complications arise when we consider the possibility that subjects who search for evidence may be concerned not with the first order question *whether p* but with some higher-order question about their attitude to *p*. For example, if higher stakes situations demand iterations of knowledge – see (Williamson, 2005) for discussion of this suggestion – then it could be entirely rational to keep searching for additional evidence after one's mind is already made up. There is some support in the empirical literature for the conjecture that self-consciousness about one's epistemic standing tends to increase epistemic anxiety; to complicate matters, increased attention to one's grasp of *p* does not always translate directly into increased accuracy with respect to *p* (Lerner & Tetlock, 1999).¹⁷ It has even been suggested that increases in cognitive energy are on the whole more directed towards putting oneself in a better position to self-justify, rather than towards increased accuracy *per se* (Mercier & Sperber, 2009), but even if increased effort heightens

¹⁶ There are further complications here, some of them arising from the possibility that someone might be seen as holding two different attitudes towards the same proposition (Friedman, MS; Petty, Tormala, Brinol, & Jarvis, 2006; Wilson, Lindsey, & Schooler, 2000). Perhaps HIGH's explicit attitude as he decides to go check the hours is one of suspension, while he has an underlying implicit attitude that the bank will be open, where the former attitude guides his current, unrushed behavior and the latter could still influence his affective states and his responses under pressure. It is not entirely clear that the underlying implicit attitude retains the full status of belief (Gendler, 2008); one might also wonder whether such a conflict in attitudes would itself tend to bar the ascription of knowledge, even if the underlying attitude were considered full belief.

¹⁷ Some work on accountability, or the expectation of having to justify oneself, is relevant here. If accountability manipulations are announced after subjects have reviewed the evidence (and presumably made up their minds), they tend to devote energy to justifying their set responses rather than rethinking them: for example, the sunk-costs error of irrational perseverance in a losing plan is exacerbated under postdecisional accountability (Conlon & Wolf, 1980) but attenuated if accountability measures are announced at the outset (Simonson & Nye, 1992). Notwithstanding these types of cases, it is thought that accountability generally enhances accuracy (Lerner 1999).

accuracy only incidentally and imperfectly, it may still end up being adaptive overall. The relationship between stakes, accountability and iterations of knowledge is in any event complex enough to warrant more detailed investigation on another occasion. For present purposes, it may be worth registering that one further way of reading the Bank case would be to see HIGH as knowing that the bank will be open while not knowing that he knows, in which case his denial of knowledge could be sincere but mistaken. Such a denial might nonetheless be conversationally appropriate, just as it could be pragmatically appropriate for him to double-check the hours, if high-stakes situations require knowing that one knows.

The central complaint against the Bank cases was that HIGH and LOW were not matched on all the attributes relevant to traditional epistemology. This type of complaint might seem to be ruled out for third-person cases which concern a single subject viewed from two angles rather than a pair of perhaps unmatched subjects in different environments. However, a related strategy is available. Whatever traditional factors it would actually be right to attribute to the evaluated subject in these third-person cases, perhaps something about the shift in perspective distorts our perception of those factors, so that the subject is from one angle or the other naturally misperceived. On this view, our intuitive capacity for knowledge ascription would always involve responsiveness to traditional factors, but the third-person cases motivating shifting-standard views would be diagnosed as supplying distorted or illusory input to this capacity.

One of the clearest such cases, Keith DeRose's Office case, succeeds in generating the appearance that a single character – Lena – might from one perspective (Thelma's) be correctly taken to know that John was at work today, while from another perspective (Louise's) being correctly taken not to know this same proposition. In the story, Lena, Thelma and Louise all share, and know they share, just the same information relevant to the proposition that John was at work today: they saw his hat in the hall, and overheard a partial conversation which seemed to involve him. In a low-stakes tavern conversation later on, Thelma seems right to describe Lena as knowing that John was at work; in a separate high-stakes conversation with the police, Louise seems right to deny that Lena knows. While these two conversations are going on, Lena herself is in some third location, not thinking about the key proposition and unaware that it might be a high stakes matter. The case stipulates that John was indeed at work.

Assuming we are satisfied that the evidence mentioned would support an evidence-based

judgment in low-anxiety conditions, the invariantist can grant that Lena knows John is at work, and that Thelma is right to say that Lena knows that John was at work. The challenge will be to explain why Louise seems right to say otherwise, and the adaptive invariantist approach will be to argue that Louise naturally feels increased epistemic anxiety when considering that proposition, and naturally imagines Lena as sharing it. As the first stage in this argument, we note that the discussion of Louise's judgment includes a first-person denial of knowledge that is reminiscent of HIGH's denial in the Bank case. Louise is first asked by the police whether she herself could testify that John was at work, and in the course of explaining why she would be unable to testify, Louise displays heightened cognitive effort – and encourages us, as readers, to feel it as well – by explicitly entertaining hypotheses under which she now realizes her evidence would be inconclusive – “John could have left his hat on the hook when he went home some previous day,” and so on (DeRose, 2009, 5). What might otherwise have been a settled question for her is re-opened under pressure, and as she wonders aloud about other possibilities, she no longer seems to have the plain outright belief that John was at work, even if she still considers this likely. The police then “ask Louise whether Lena might know that John was in,” and Louise responds: “She has the same reasons I have for thinking John was there, but like me, she doesn't know that John was there.” (ibid.)

If Louise's own epistemic anxiety on the question of John's whereabouts is enough to bar her from reaching an evidence-based outright belief on that question, this anxiety would have an impact on the way she would naturally simulate or reconstruct Lena's state of mind. If Louise now finds her own evidence inadequate, and if contemplation of her stakes induces a similar frame of mind in us as readers, it will be hard for her (and for us) to imagine Lena's rational satisfaction with the same evidence. As far as the proposition about John is concerned, the scenario makes Lena's position more naïve than either the reader's or Louise's – together with Louise, we are thinking harder about the evidence, and actively entertaining various hypothetical scenarios relevant to the key proposition, while Lena is oblivious. We have a well-documented tendency to misread the mental states of those who are more naïve than we are, to evaluate them as though they were privy to our concerns, without being aware that we are doing so (Birch, 2004; Birch & Bloom, 2007; Nickerson, 1999). There is some reason to think that if we are epistemically anxious, we will automatically tend to evaluate Lena as if she shared our heightened need for evidence (Nagel, 2010).¹⁸ Alternatively, perhaps what the police do in asking Louise whether Lena might know that

¹⁸ Objection: to the extent that we overshare our privileged information, why wouldn't we also project onto Lena our knowledge that John was actually at work? When we are evaluating Lena's state of mind with

John was at work is to make Louise think about what the police would find if they were to go question Lena themselves. Having found that the pressure to testify has wilted her own belief about John's whereabouts, Louise naturally imagines that Lena will suffer the same experience under pressure. Louise could be giving the police a literally correct report of the matter that interests them; namely, the state of mind that they would encounter in Lena if they question her.

There may be some good reason to understand the police's "might Lena know?" question as "will she know, if we go and ask?" rather than "does she know, right now, as she sits alone on the sofa thinking about other things?" Our capacities for intuitive mental state ascription are most naturally deployed in explaining and predicting the assertions and actions of others; it is arguably easier to judge someone to think or know some proposition if we imagine them saying it or acting upon it.¹⁹ When we imagine people as possible informants, we naturally imagine them in cooperative conversation with us, in a position where they would come to share our heightened concerns. Because the police are asking Louise about Lena's capacity to testify, it will be natural for her to imagine Lena in that most salient position, and for us, as readers of the case, to follow in Louise's footsteps as we evaluate what Louise thinks about what Lena would say. If we apply considerable effort to focusing on the difference between Louise's stressed position and Lena's relaxed position, it no longer seems so clear that what Louise is saying about Lena is correct.²⁰

Adaptive invariantism agrees with IRI that the subject who knows her stakes are high will need more evidence to gain knowledge than her low-stakes counterpart. Views committed to this principle are thought to have difficulty with counterfactuals of the form "He knows, but he wouldn't have known if more had been at stake" (DeRose, 2009, 194). If we intuitively require greater

respect to this very proposition, we need to represent her as making some transition from her resources to this target. We may naturally misrepresent these resources, but it is unlikely that we would do so by including among them the stipulation about the target proposition that we have been given in our "God's eye" view of the case: a stipulation of the truth of the proposition cannot be represented as included in her resources for making a judgment on that very proposition. Even if our knowledge of the stipulation would incline us to see her as accepting this proposition, we would need to bracket her acceptance of the target proposition in evaluating her reasoning, just as we need to bracket our own acceptance of a proposition when we raise the question of whether we really know it to be true (Williamson, 2008). I am grateful to John MacFarlane for pressing me on this point.

¹⁹ In fact, imagining or witnessing ourselves or others speaking or acting may be indispensable to all propositional attitude ascription (on this point see Carruthers, 2009).

²⁰ DeRose himself observes that intuitions are cloudier if we think about Thelma's relaxed position while evaluating Louise's claim, or the converse (DeRose, 2009, 5). Most context-sensitive terms do not seem to provoke this kind of confusion: in accepting simultaneous utterances of "I am hungry" and "I am not hungry" where there are two speakers, we do not experience any drop in the acceptability of either utterance if we think of the other at the same time.

evidence from a higher-stakes subject, why do these counterfactuals sound jarring? One possible source of resistance is that we ordinarily expect an ordinary rational subject who is shifted into higher stakes to feel increased epistemic anxiety and do the evidence-collecting work appropriate to his increased anxiety. If I see a subject as knowing that p , I see him as being appropriately eager for new evidence on the question of whether p ; I ordinarily expect him to keep knowing p (even if this means securing more evidence) if it becomes a high-stake matter. Of course, careful stipulations in the case could make it clear that no further evidence will be available. However, once we have emphasized the restrictions on evidence and the prospect of rising stakes, it will be natural for readers to misrepresent even the casual subject of the antecedent of the conditional as though he already had the anxieties he would suffer in the consequent.

A similar strategy can be used to explain reactions to Ignorant High Stakes cases, where a subject has enough evidence to support a Low Stakes judgment, but is in fact in a High Stakes predicament without knowing that she is. In at least some cases, we may become more stringent in ascribing knowledge to subjects who are described as having much more at stake than they know (Stanley, 2005). As readers of the case whose attention is drawn to the high stakes situation, we could be expected to feel heightened epistemic anxiety with respect to the key proposition, in a way that could obstruct us from appreciating the subject's low-anxiety reasoning as acceptable (Nagel, 2008). Any sense that the Ignorant High Stakes subject fails to know would be generated by a natural misrepresentation of this subject's mental state; negative intuitions about Ignorant High Stakes subjects could then be dismissed as illusory. Unlike the advocate of IRI, the adaptive invariantist does not take a subject's actual stakes to be a factor in whether she has knowledge; stakes have at most an indirect impact, mediated by their influence on traditional factors. Because subjects in high-stakes contexts are often aware of their stakes, and because awareness of high stakes lifts epistemic anxiety and produces the need for greater evidence, a shift in stakes will often change what the subject needs to do in order to comply with the traditional norm of thinking in an evidence-based manner. But the standard that matters to evaluating the subject can most simply be cast in terms of the traditional norm, with no direct reference to stakes. Because the Ignorant High Stakes subject appears to violate this norm, he is intuitively represented as failing to know; because he does not actually violate the norm, he knows as well as his Low Stakes counterpart.

Here an advocate of IRI might object that we could be judging the Ignorant High Stakes subject in precisely the right way, against the benchmark of the level of anxiety he *ought* to be feeling. To the extent that the subject is in fact in a high-stakes predicament, would it not be in some

sense appropriate for him to be feeling increased epistemic anxiety? The IRI-friendly claim that low anxiety would in High Stakes situations be a bar to knowledge could be supported in part by a comparison to another type of case in which the subject's epistemic anxiety is arguably insufficient for his circumstances – namely, the Gettier case. The unsuspecting Gettier victim who sets his watch on a broken clock is making a routine, low-anxiety evidence-based judgment, as we can appreciate when we evaluate his judgment as justified. If he knew about the state of the clock he'd be more anxious, and seek further evidence before making up his mind about the time. Both Ignorant High Stakes and the Gettier victim can be seen from our more enlightened point of view to be insufficiently epistemically anxious, and if both types of subject are intuitively judged to lack knowledge, one might think that any arguments to dismiss the first type of intuition as illusory should equally undermine the second.

To respond to this objection, the adaptive invariantist who wants to maintain the traditional view that Gettier victims lack knowledge can search for relevant differences between Ignorant High Stakes and Gettier cases. Granting that both of them would feel more anxious if they knew what we did, and in some sense should feel more anxious than they are feeling, the adaptive invariantist can observe that the relevant sense of “should” is not quite the same for the two types of naïve subject. The underlying reason why we should feel increased epistemic anxiety in high-stakes circumstances has to do with what is pragmatically desirable: the pragmatic benefits of increased accuracy offset the pragmatic costs of increased effort and time spent searching for evidence. If Ignorant High Stakes does not feel high anxiety in her current predicament, he is at risk of failing to maximize utility. The standard Gettier victim, on the other hand, should be collecting more evidence because she is in a position where her belief forming methods have a higher risk of error.²¹ We might get the sense that Ignorant High Stakes is unreliable if we benchmark his performance against what we would expect from someone like him who knew his stakes; meanwhile the Gettier victim actually is objectively unreliable, or at risk of error, given her circumstances (which is not, of course, to suggest that she is epistemically blameworthy). The advocate of IRI may protest that it is question-begging to separate pragmatic and epistemic considerations in this manner, but the burden of argument probably lies on those who want to

²¹ There are some nonstandard types of Gettier case, and arguably not all of them involve a risk of unreliability. Harman's dead dictator case, for example, raises worries about the subject's risk of being undermined by misleading defeaters. Here the traditional norm that matters may have to do with the durability or resilience of a belief rather than its accuracy. A full treatment of the relationship between epistemic anxiety and Gettier cases is a project for another occasion.

undermine this traditional distinction. As long as we find it intuitively unpalatable to, say, accept as knowledge those wishful thinking cases in which the anxious subject curtails evidence collection to maximize utility, we should have some qualms about the suggestion that knowledge requires our having levels of epistemic anxiety that are set to maximize pragmatic value. If the cases motivating IRI can be explained by appeal to the traditional factors of reliability and confidence, and to known biases in mental state ascription, then considerations of simplicity should discourage us from counting practical interests as a further knowledge-determining factor.

If the adaptive invariantist is right to maintain that stakes themselves only matter indirectly, she should be able to sidestep one potential problem having to do with the calibration of our responses to rising and falling stakes. Section 1 presented evidence that we are only roughly adaptive in our thinking: it is an open question precisely how well we balance increases in effort against increases in accuracy. We automatically increase effort under rising stakes, but it may well be the case that these automatic processes are insufficient to maximize utility; perhaps we, say, naturally waste effort by thinking a bit too hard when stakes get higher. However, as long as the automatic processes underpinning approximate adaptivity are roughly the same in all of us, imperfect calibration may not matter to the quality of our epistemic assessments.²² If I expect subjects to exhibit certain patterns of cognitive effort as stakes rise and fall, and if subjects are in general governed by the same rules of cognitive effort, then whether or not those rules are in fact pragmatically optimal, it will be right to worry about any subject who deviates downwards from the patterns of evidence-seeking naturally generated by those rules.²³ If she stops collecting evidence before attaining the threshold of evidence-collection that would be natural for her predicament, I would be right to worry that she is distracted or thinking wishfully, even if this pattern of evidence-collection would in fact maximize her utility.

There are much deeper worries nearby. This paper has presented adaptive invariantism strictly as a theory of intuitive knowledge ascription; it has not presented a theory of knowledge. The aim has been to explain various patterns of positive and negative intuitions, in part by drawing

²² There is evidence of some individual variation in calibration, but it is not clear that this is any worse than the individual variation in visual acuity, or any other cognitive resource: those of us with poor vision may mildly but systematically underestimate what others can see, but this familiar problem grounds caution rather than skepticism about the epistemic assessments affected by this individual limitation, and supports our general practice of checking epistemic assessments with others.

²³ There could be conditions under which subjects deviate upwards from what we naturally expect – double-checking under low stakes, and so forth. We have a well-established general tendency to minimize cognitive effort, or to be “cognitive misers”, in Stanovich’s expression (Stanovich & West, 2000, 2008), but there is individual variation here, and there are also conditions under which we can be stimulated to find cognitive effort unusually satisfying (Kruglanski & Webster, 1996). Of course we would not see confident subjects as failing to know where they have done *more* than is naturally expected.

attention to the role played by the natural processes automatically regulating cognitive effort. Describing those natural processes leaves it somewhat open how we should evaluate them. A low-stakes subject who makes up his mind on moderate evidence may intuitively register for us as a knower, but we can wonder whether this intuitive assessment should be taken at face value, whether such a subject really knows or merely seems to know. It is one thing to understand how we instinctively measure others against our own default ways of thinking, and quite another to examine the merits of those default ways of thinking. In examining those merits we face many new challenges lying beyond the scope of the present work. But one of the challenges we face is by now familiar: when we wonder whether our low-stakes subject is thinking as he ought, we will naturally consider other ways in which he could be thinking, further pieces of evidence he could collect, further hypotheses he could consider. Just by thinking about the possibility of him trying harder, we ourselves come to think harder, increasing our own epistemic anxiety on the proposition of interest, and increasing our own propensity to find his thinking unsatisfactory. Direct attention to the question of how much epistemic anxiety to apply to a given question is a good way of increasing that anxiety, unless we can tame that impulse, perhaps by finding more oblique ways to think about the problem. Learning more about the natural mechanisms governing cognitive effort might be one way to brace ourselves against their risky side-effects. This is one of the ways in which the development of a descriptively accurate theory of intuitive knowledge ascription could contribute to the eventual development of a theory of knowledge.²⁴

²⁴ Thanks to Sergio Tenenbaum and Jonathan Weisberg for comments and discussion. Thanks also to the Social Sciences and Humanities Research Council of Canada for supporting this research.

References

- Alter, A., & Oppenheimer, D. (2009). Uniting the tribes of fluency to form a metacognitive nation. *Personality and Social Psychology Review*, 13(3), 219.
- Alter, A., Oppenheimer, D., Epley, N., & Eyre, R. (2007). Overcoming intuition: Metacognitive difficulty activates analytic reasoning. *Journal of Experimental Psychology: General*, 136(4), 569-576.
- Bach, K. (2005). The Emperor's New 'Knows'. In G. Preyer & G. Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (pp. 51-89). New York: Oxford University Press.
- Birch. (2004). Understanding children's and adults' limitations in mental state reasoning, *Trends in cognitive sciences* (Vol. 8, pp. 255).
- Birch, S., & Bloom, P. (2007). The curse of knowledge in reasoning about false beliefs. *Psychological Science*, 18(5), 382.
- Broder, A., & Newell, B. (2008). Challenging some common beliefs: Empirical work within the adaptive toolbox metaphor. *Judgment and Decision Making*, 3(3), 205-214.
- Carruthers, P. (2009). How we know our own minds: the relationship between mindreading and metacognition. *Behavioral and Brain Sciences*, 32(02), 121-138.
- Cohen, S. (1999). Contextualism, Skepticism, and the Structure of Reasons (Volume 13: Epistemology). *Nous-Supplement: Philosophical Perspectives*, 13, 57-89.
- Conlon, E., & Wolf, G. (1980). The moderating effects of strategy, visibility, and involvement on allocation behavior: An extension of Staw's escalation paradigm. *Organizational behavior and human performance*, 26(2), 172-192.
- Connolly, T., & Serre, P. (1984). Information search in judgment tasks: The effects of unequal cue validity and cost. *Organizational behavior and human performance*, 34(3), 387-401.
- Davies, M., & Gardner, D. (2010). *Frequency Dictionary of American English*. New York: Routledge.
- DeRose, K. (1992). Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52(4), 913-929.
- DeRose, K. (2009). *The Case for Contextualism: Knowledge, Skepticism, and Context, Volume 1*. New York: Oxford University Press.
- Eagly, A., & Chaiken, S. (1993). *The psychology of attitudes*: Harcourt Brace Jovanovich Fort Worth, TX.
- Epley, N., & Gilovich, T. (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(3), 199-212.
- Evans, J. (2007). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, 59, 255-278.
- Evans, J. S. B. T. (2006). The heuristic-analytic theory of reasoning: Extension and evaluation. *Psychonomic Bulletin & Review*, 13(3), 378 - 395.
- Fantl, J., & McGrath, M. (2009). *Knowledge in an uncertain world*: Oxford: Oxford University Press.
- Foley, R. (1993). *Working without a Net*. New York: Oxford University Press.
- Foley, R. (2009). Beliefs, degrees of belief, and the Lockean thesis. In F. Huber & C. Schmidt-Petri (Eds.), *Degrees of Belief* (pp. 37-47). Konstanz: Springer.
- Frankish, K. (2009). Systems and levels: dual-system theories and the personal-subpersonal distinction. In *Two Minds: Dual Processes and Beyond*, 89-107.
- Frankish, K., & Evans, J. (2009). The duality of mind: an historical perspective. In K. Frankish & J. Evans (Eds.), *In Two Minds: Dual Process Theory and Beyond* (pp. 1-29). Oxford: Oxford University Press.
- Friedman, J. (MS). *Suspended Judgment*. Unpublished manuscript.
- Ganson, D. (2008). Evidentialism and pragmatic constraints on outright belief. *Philosophical Studies*, 139(3), 441-458.
- Gendler, T. (2008). Alief and belief. *Journal of Philosophy*, 105(10), 634-663.
- Gigerenzer, G., & Todd, P. (1999). *Simple heuristics that make us smart*: Oxford University Press, USA.
- Goldman, A. (2006). *Simulating minds: The philosophy, psychology, and neuroscience of mindreading*. New York: Oxford University Press.
- Gollwitzer, P., & Bayer, U. (1999). Deliberative versus implemental mindsets in the control of action. *Dual-process theories in social psychology*, 403-422.
- Gopnik, A., & Meltzoff, A. (1997). *Words, Thoughts, and Theories*. Cambridge, MA: MIT Press.
- Harman, G. (1986). *Change in view: Principles of reasoning*. Cambridge, MA: MIT Press.

- Hart, J. T. (1965). Memory and the feeling-of-knowing experience. *Journal of Educational Psychology*, 56(4), 208-216.
- Hausmann, D., & Lage, D. (2008). Sequential evidence accumulation in decision making: The individual desired level of confidence can explain the extent of information acquisition. *Judgment and Decision Making*, 3(3), 229-243.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. New York: Oxford University Press.
- Hawthorne, J. (2005). The case for closure. In M. Steup & E. Sosa (Eds.), *Contemporary debates in epistemology* (pp. 26-43). New York: Blackwell.
- Hookway, C. J. (1998). Doubt: Affective States and the Regulation of Inquiry. In C. Misak (Ed.), *Pragmatism*. Calgary: University of Calgary Press.
- Hulland, J., & Kleinmuntz, D. (1994). Factors influencing the use of internal summary evaluations versus external information in choice. *Journal of Behavioral Decision Making*, 7(2), 79-102.
- Kelley, C. (1993). Remembering mistaken for knowing: Ease of retrieval as a basis for confidence in answers to general knowledge questions. *Journal of Memory and Language* (Vol. 32, pp. 1).
- Kelley, H., Jones, E., Kanouse, D., Nisbett, R., Valins, S., & Weiner, B. (1972). Attribution in Social Interaction. In Anonymous (Ed.), *Attribution: Perceiving the Causes of Behavior*. Morristown, NJ: General Learning Press.
- Koriat, A. (1993). How do we know that we know? The accessibility model of the feeling of knowing. *Psychological review*, 100(4), 609-639.
- Koriat, A. (1995). Dissociating knowing and the feeling of knowing: Further evidence for the accessibility model. *Journal of Experimental Psychology: General*, 124(3), 311-333.
- Koriat, A. (2000). *Control processes in remembering*. New York, NY, US: Oxford University Press.
- Koriat, A., & Goldsmith, M. (1996). Monitoring and control processes in the strategic regulation of memory accuracy. *Psychological review*, 103(3), 490-517.
- Koriat, A., & Levy-Sadot, R. (2000). Conscious and Unconscious Metacognition: A Rejoinder. *Consciousness and Cognition*, 9(2), 193-202.
- Kruglanski, A. W., & Webster, D. M. (1996). Motivated closing of the mind: "Seizing" and "Freezing". *Psychological review*, 103(2), 263-283.
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3), 480-498.
- Lee, M., & Cummins, T. (2004). Evidence accumulation in decision making: Unifying the take the best and the rational models. *Psychonomic Bulletin & Review*, 11(2), 343-352.
- Lerner, J. S., & Tetlock, P. E. (1999). Accounting for the Effects of Accountability. *Psychological Bulletin*, 125(2), 255-275.
- Mayselless, O., & Kruglanski, A. (1987). What makes you so sure? Effects of epistemic motivations on judgmental confidence. *Organizational Behavior and Human Decision Processes*, 39, 162-183.
- McAllister, D. W., Mitchell, T. R., & Beach, L. R. (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(2), 228-244.
- Mercier, H., & Sperber, D. (2009). Intuitive and reflective inferences. *J. St. BT Evans & K. Frankish (Eds.), In two minds: Dual processes and beyond*. Oxford, UK: Oxford University Press.
- Nagel, J. (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*, 86(2), 279-294.
- Nagel, J. (2010). Knowledge ascriptions and the psychological consequences of thinking about error. *Philosophical Quarterly*, 60(239), 286-306.
- Nagel, J. (forthcoming). The psychological basis of the Harman-Vogel paradox. *Philosophers' Imprint*.
- Newell, B. (2005). Re-visions of rationality? *Trends in cognitive sciences*, 9(1), 11-15.
- Newell, B., & Lee, M. (2010). The Right Tool for the Job? Evidence Accumulation in Decision Making. *Psychological review*.
- Newell, B., Rakow, T., Weston, N., & Shanks, D. (2004). Search strategies in decision making: The success of 'success'. *Journal of Behavioral Decision Making*, 17(2), 117-137.
- Newell, B., & Shanks, D. (2003). Take the best or look at the rest? Factors influencing 'one-reason' decision making. *Journal of Experimental Psychology-Learning Memory and Cognition*, 29(1), 53-65.
- Newell, B., Weston, N., & Shanks, D. (2003). Empirical tests of a fast-and-frugal heuristic: Not everyone 'takes-the-best'. *Organizational Behavior and Human Decision Processes*, 91(1), 82-96.

- Nickerson, R. S. (1999). How we know--and sometimes misjudge--what others know: Imputing one's own knowledge to others. *Psychological Bulletin*, 125(6), 737-759.
- Payne, J., Bettman, J., & Johnson, E. (1993). *The Adaptive Decision Maker*. New York: Cambridge University Press.
- Perner, J. (1991). *Understanding the Representational Mind*: MIT Press Cambridge, MA.
- Petty, R., Tormala, Z., Brinol, P., & Jarvis, W. (2006). Implicit ambivalence from attitude change: An exploration of the PAST model. *Journal of Personality and Social Psychology*, 90(1), 21.
- Pinillos, A. (submitted). Knowledge, experiments, and practical interests.
- Rieskamp, J., & Otto, P. (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology: General*, 135, 207-236.
- Rysiew, P. (forthcoming). Surveys, Intuitions, Knowledge Attributions: Comments on Keith DeRose's "Contextualism, Contrastivism, and X-Phi Surveys". *Philosophical Studies*.
- Schaffer, J., & Knobe, J. (forthcoming). Contrastive Knowledge Surveyed. *Nous*
- Schwarz, N. (1998). Accessible content and accessibility experiences: The interplay of declarative and experiential information in judgment. *Personality and Social Psychology Review*, 2(2), 87.
- Shah, A., & Oppenheimer, D. (2008). Heuristics made easy: An effort-reduction framework. *Psychological bulletin*, 134(2), 207-222.
- Simonson, I., & Nye, P. (1992). The effect of accountability on susceptibility to decision errors. *Organizational behavior and human decision processes*, 51(3), 416-446.
- Sloman, S. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119(1), 3-22.
- Stanley, J. (2005). *Knowledge and Practical Interests*. New York: Oxford University Press.
- Stanovich, K., & West, R. (2000). Individual differences in reasoning: implications for the rationality debate? *The Behavioral and brain sciences*, 23(5), 645.
- Stanovich, K., & West, R. (2008). On the relative independence of thinking biases and cognitive ability. *Journal of personality and social psychology*, 94(4), 672.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. Roth & G. Ross (Eds.), *Doubting: Contemporary Perspectives on Skepticism* (pp. 13-27). Dordrecht: Kluwer.
- Weatherston, B. (2005). Can we do without Pragmatic Encroachment? *Philosophical Perspectives*, 19(1), 417-443.
- Williamson, T. (2005). Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly*, 55, 213-235.
- Williamson, T. (2008). Knowledge and Scepticism. In F. Jackson & M. Smith (Eds.), *Oxford Handbook of Contemporary Philosophy* (pp. 681-700). New York: Oxford University Press.
- Wilson, T., Lindsey, S., & Schooler, T. (2000). A model of dual attitudes. *Psychological Review*, 107(1), 101-126.