

Are Knowledge Ascriptions Sensitive to Social Context?

Alexander Jackson

Abstract: Plausibly, how much is at stake in some salient practical task can affect how generously people ascribe knowledge of task-relevant facts. There is a metaphysical puzzle about this phenomenon, and an empirical puzzle. Metaphysically: there are competing theories about when and how practical stakes affect whether it is correct to ascribe knowledge. Which of these theories is the right one? Empirically: experimental philosophy has struggled to find a stakes-effect on people's knowledge ascriptions. Is the alleged phenomenon just a philosopher's fantasy? I propose a new psychological account of when and why people's knowledge ascriptions are sensitive to stakes. My hypothesis is motivated by empirical research on how people's judgements are sensitive to their social context. Specifically, people's evaluations are sensitive to their 'psychological distance' from the scenarios they are considering. When using 'fixed-evidence probes', experimental philosophy has found that what's at stake for a fictional character in a made-up scenario has little or no effect on how participants ascribe knowledge to them. My hypothesis predicts this finding: the scenarios are too 'psychologically distant' to participants. Our empirical puzzle is resolved: the stakes-effect often present in the wild won't be present in vignette studies. (This illustrates a widespread problem with X-phi vignette studies: if people might judge differently in other social contexts, we can't generalize from the results of these experiments. That is, vignette studies are of doubtful 'external validity'.) The hypothesis also resolves our metaphysical puzzle. It predicts that people do not ascribe knowledge in a way deemed correct by any of the standard philosophical views, namely classical invariantism, interest-relative invariantism, and contextualism. Our knowledge ascriptions shift around in the way that's most useful for social beings like us, and this pattern in our judgements can only be endorsed by a genuinely relativist metaphysics for knowledge.

Keywords knowledge; pragmatic encroachment; epistemic contextualism; epistemic relativism; experimental philosophy; social cognition.

1. Intuitions and metaphysics, X-phi and psychology.

1.1 *The philosophical debate about bank cases.*

Philosophers debate what metaphysical or semantic moral to draw from the following sort of example—a ‘bank case’.¹

It’s Friday afternoon, and there’s a long line at the bank. Thus Bill decides to leave depositing his cheque till Saturday. It’s not important that he deposit it immediately. Bill is stopped in the street by Hannah, who asks him whether the bank will be open on Saturday. He answers, “The bank will be open on Saturday. I know because I was there on a Saturday a couple of months ago.” It is really important to Hannah and Sarah that they deposit their cheque before Monday—they have a big direct debit due. They also remember that the bank was open on a Saturday two months ago, but they want to be sure about this Saturday. Hannah reports back to Sarah, “No, that guy doesn’t know that the bank will be open on Saturday.”

Philosophers have been puzzled by their intuitions as to whether Bill knows the bank will be open on Saturday. At the start of the vignette, when Bill’s low-stakes decision is in view, we think he does know. But at the end of the vignette, when Hannah and Sarah’s high-stakes choice is salient, we agree with Hannah that Bill doesn’t know. It seems that we issue contrary verdicts as to whether Bill knows, and thus that one of those verdicts is mistaken. Which verdict is in error, and why are we lead astray? Or is there a way to reconcile the apparently conflicting intuitions as to whether Bill knows? Let’s review three leading answers.

Interest-Relative Invariantism (IRI) says that Bill knows that the bank will be open, but Hannah and Sarah can’t know it even though they have the same evidence as Bill. That’s so, according to IRI, because whether someone knows something depends partly on how bad it would be for them have a false belief on the matter. The stakes are a lot higher for

¹ DeRose (1992; 2009: 1–9) and Stanley (2005: 3–5) give cases like this; Cohen’s (1999) ‘airport’ case is similar. The example should not also vary which possibilities of error are salient.

Hannah and Sarah than they are for Bill, and so the evidence is strong enough for Bill to know that the bank will be open, but not for Hannah to do so. On this view, Bill speaks the truth when he says that he knows the bank will be open, and Hannah speaks falsely when she denies that he knows. We go wrong when we come to the end of the vignette and assess whether Bill knows using the standard that's appropriate to Hannah's practical situation, not his own.²

Contextualism says that the reference of the word “knows” varies, depending on the speaker's context. Contextualism *per se* is not committed to any particular account of bank cases, but we'll focus on versions that do give a distinctive treatment. Contextualists can say that Bill spoke the truth when he said he “knows”, and Hannah speaks the truth when she says Bill doesn't “know”, as the two speakers are talking about different things. At a first pass, Bill semantically expresses that he knows-by-lax-standards. Those are the standards it is relevant for him to be talking about, given that it wouldn't be a disaster for him to falsely believe that the bank will be open on Saturday morning. But Hannah semantically expresses that Bill does not know-by-stringent-standards. Those are the standards it is relevant for her to be talking about, given that it would be a disaster for her to falsely believe that the bank will be open. Bill knows-by-lax-standards that the bank will be open, but does not know-by-stringent-standards; so Bill and Hannah both speak the truth. On this view, our mistake is to think that we have changed our minds over the course of the vignette about some question, namely whether Bill knows *tout court*. There is no such question. Rather, we have compatible intuitions about different questions.³

Classical invariantism rejects both IRI and contextualism. According to classical invariantism, Bill asserts that he knows and Hannah denies it, contra contextualism. But the strength of evidence needed to know something is the same for all subjects, regardless

² Interest-relative invariantism is defended by Hawthorne (2004 chapter 4—tentatively), Stanley (2005), Fantl & McGrath (2009, 2012), and Weatherson (2012).

³ Contextualists include DeRose (1992, 2009), Cohen (1999), Lewis (1996), Greco (2008: 432–5), Fricker (2008), Henderson (2011), Hannon (2013), McKenna (2013), and Blome-Tillmann (2014). ‘WAM-ing invariantism’ is the view that while “knows” always semantically expresses the same thing, people use the word to conversationally implicate a variety of other epistemic states, with results similar to contextualism; see §6.

of what's at stake for them, contra IRI. Either Bill and Hannah are both in an epistemic position to know that the bank will be open, or neither is. According to classical invariantism, one of our verdicts about whether Bill knows is simply wrong. Different accounts of why we are misled are possible.⁴

Until recently, all sides agreed on the pattern of intuitions the bank case evokes: first it seems that Bill spoke the truth in saying that he “knows”, and then it seems that Hannah spoke the truth in saying he doesn't “know”. Philosophers agreed that the practical stakes affect our intuitive knowledge ascriptions. They disagreed about which of those intuitions should be rejected, if any. However, findings from experimental philosophy ('X-phi') raise the spectre that this debate is based on a mirage. It seems that ordinary people don't share philosophers' intuitions about the bank case and suchlike.

1.2 Perplexing results from experimental philosophy.

Standard studies use a 'fixed-evidence probe'. Two versions of a vignette are prepared, one in which the character is in a high-stakes practical situation (like Hannah), and the other in which the character is in a low-stakes situation (like Bill). Everything else is held fixed between the two versions, including the evidence the character has that bears on some practically relevant proposition (such as that the bank will be open on Saturday morning). Participants read one version, and indicate the degree to which they agree that the character knows the practically relevant proposition. We compare the spread of judgements from the low-stakes version of the vignette with those from the high-stakes version. This tells us whether the practical stakes for the character affects participants' judgements as to whether the character knows.

These studies don't find the expected large effect of practical stakes on people's knowledge ascriptions (see §7). A few studies found statistically significant evidence of a small effect, but more—including a huge cross-cultural study (Rose et al. 2019)—found no evidence of an effect. I will generally summarize these findings by saying that studies using fixed-evidence probes find no effect of stakes on participants' knowledge ascriptions.

⁴ Classical invariantism is defended by Williamson (2005), Nagel (2010), Gerken (2017 chapter 12), and Dinges (2019 chapter 8).

This seems to me a reasonable moral to draw at present. Maybe things are more complex and there is a small stakes effect; I address this possibility in §8.

If we focus on results of testing bank cases and alike, it seems that ordinary people do not share philosophers' intuitions that generate the philosophical debate. Some might conclude that bank cases are philosophically unproblematic—there aren't putatively contradictory intuitions needing to be reconciled. But I think the empirical findings deepen the puzzle, for three reasons.

First, another trend in the X-phi results complicates the empirical picture. Studies that use an 'evidence-seeking probe' ask participants questions like, "How many times does the student need to proof-read their essay before they know it contains no typos?" Again, half the participants read a high-stakes version of the vignette, and half a low-stakes version. These studies tend to find an effect of stakes on participants' answers: people write down bigger numbers when the stakes are higher (see §8). Why is there apparently a stakes effect on responses to evidence-seeking probes, and not on responses to fixed-evidence probes?

Second, the philosophical dispute arose because many philosophers found their intuitions about cases *are* sensitive to what was at stake for the character. Why do philosophers differ in this respect from typical participants in fixed-evidence surveys?

Third and most importantly, it remains terrifically plausible that ordinary people's knowledge ascriptions are stakes-sensitive in the wild when dealing with matters they care about, rather than imagining predicaments for unfamiliar fictional characters. So I shall argue below (and so says every psychologist with whom I've discussed the matter). Why don't vignette studies elicit the stakes-sensitivity still plausibly present in the wild? More generally, *when* are people's knowledge ascriptions sensitive to the stakes in some salient practical task? Currently, we have data about people's responses to vignette surveys, but not from ecologically valid experiments that approximate real-life situations. It would be mistaken to put theorizing on hold until such experiments have been performed. Experiments should be designed to test sensible hypotheses. I will construct a well-motivated hypothesis for future empirical work to test.

Our hypothesis should predict the findings of experimental philosophy, and make plausible predictions about knowledge ascriptions in the wild. There is another empirical

constraint on our hypothesis: it should be guided by well-supported theories about how the mind works. In particular, our hypothesis about knowledge-ascriptions should cohere with general theories of evaluation and belief-formation. The standard psychological theories reviewed in §§2–3 powerfully constrain which hypotheses are plausible.

1.3 Outline of the paper.

In my view, practical stakes *sometimes* affect people’s knowledge ascriptions. I will present a hypothesis about how and when. My hypothesis is motivated by well-supported theories in psychology, reviewed in §§2–3. In particular, empirical research finds that people’s evaluations are sensitive to their ‘psychological distance’ from the scenarios they are considering. I propose that people are more likely to adopt an epistemic standard appropriate to a given practical task when that task is ‘psychologically close’ to them—when it is a task in the here and now for someone socially close. I formulate the psychological hypothesis in §4, and draw out its predictions in §§5–7.

For one thing, the hypothesis predicts that participants in fixed-evidence surveys won’t be swayed by what’s at stake for the fictional characters they read about (§7). Participants are too ‘psychologically distant’ from the fictional characters described in the vignettes. The hypothesis accurately predicts that fixed-evidence X-phi surveys won’t find a stakes effect. Yet it also predicts that knowledge ascriptions are often stakes-sensitive in the wild. This illustrates a general problem with the vignette surveys typical of current experimental philosophy: when social context might influence cognition, one cannot use vignette survey results to draw conclusions about how people think in real life.

§8 accounts for two extra phenomena that contrast with the findings of fixed-evidence surveys, namely philosophers’ intuitions and the results from evidence-seeking probes. My explanations appeal to factors other than psychological distance. The hypothesis of §4 is a *ceteris paribus* generalization, and §8 treats the two extra phenomena as exceptions to the default way of ascribing knowledge.

What light does the psychological hypothesis shed on the metaphysical and semantic dispute about knowledge ascriptions? The hypothesis predicts that people do not ascribe knowledge in a way deemed correct by any of the standard philosophical views, namely classical invariantism, interest-relative invariantism, and contextualism (§§5–7).

Throughout the psychological discussion, I emphasize that the knowledge ascriptions the hypothesis predicts are the practically useful evaluations for those judges to make. In §9, I address what metaphysical and semantic conclusions we should draw, assuming my psychological hypothesis is correct. What philosophical view of knowledge should we take?—Not one of the standard ones. We should legitimize people’s usefully shifting knowledge ascriptions, and that requires a genuinely relativist metaphysics—maybe truth-relativism, maybe something else, but something bold.

2. Psychological background: occurrent full belief can vary with the stakes.

How much is at stake in the task at hand can affect whether someone will form an occurrent full belief in a proposition. That’s the premise this section articulates. For example, I’d typically judge that my car won’t be stolen if I leave it on Pine Street for two hours. However, if there were a million dollars cash in the glove box, then I’d merely assign a high probability to the prediction that the car would not be stolen, and consider alternatives that would reduce the risk, such as parking in a garage or reducing the time the car is left unattended. This hypothesis is pre-theoretically appealing. Moreover, it is well-motivated on scientific grounds, as I shall now explain (building on Nagel 2010 and Gerken 2017: 263–5).

The hypothesis distinguishes two types of mental state: outright judgements about how things are, and assignments of subjective probabilities. We can refer to these states as occurrent full beliefs and occurrent partial beliefs respectively, where an occurrent mental state is an active mental state (Bartlett 2018). By holding an occurrent full belief, one treats the matter as settled; not so if one holds an occurrent partial belief. Full or outright belief is the mental state naturally expressed by assertion. Ordinary talk of ‘belief’ sometimes refers to tentative belief, a different mental state in which one does not take it to be settled that p.⁵

⁵ Philosophers endorsing this notion of outright judgement include Harman (1986, chapters 3 & 5), Nagel (2010), Weisberg (2016), Holton (2014), Staffel (2013, 2019) and Friedman (2019).

Occurrent partial belief is very different from ‘credence’, a dispositional notion defined in terms of betting behavior employed by decision theorists and formal epistemologists (Easwaran 2011a, 2011b). Nor should occurrent partial belief be tied to the ‘epistemic’ or ‘evidential’ probabilities

There is a compelling explanation for why we make outright judgements, rather than only assigning probabilities: outright judgements make decision making computationally tractable (Harman 1986, chapters 3 & 5; Staffel 2019). If one assesses a possible course of action based on four full beliefs about how things will transpire, one need only consider the single future in which those premises are all true. For example: there will be a game of cricket played at Lord's tomorrow, the trains will run as advertised, it won't rain, and nice lunches will be available; that prospect is appealing. But if one has four occurrent partial beliefs in the relevant propositions instead, then one has sixteen possible futures to consider (or to selectively ignore). Evaluating sixteen possible futures is prohibitively slow and computationally expensive. This is especially so if one simulates possible futures and then evaluates them (Gilbert & Wilson 2007, Kahneman 2003). This can only be done one future at a time, as it monopolizes working memory (Evans & Stanovich 2013: 235–6). Relying on occurrent full beliefs is the most computationally frugal way of simplifying the decision problem. It eliminates 15 of the possibilities from consideration without calculating each of their probabilities.

Given that the function of outright judgement is to simplify decision making, we should expect people's perception of how much is at stake to affect whether they'll form a judgement given certain evidence. It is computationally easier to make a decision on the basis of an occurrent full belief, such as that my car won't be stolen if I leave it on Pine Street for two hours, than on the basis of assigning the prediction a high probability. However, there are situations in which it is worth paying the computational cost so as to gain a more nuanced conception of the situation. If there is a million dollars cash in the glove box, then it is better to assign a very high probability to the prediction that the car would not be stolen, and consider alternatives that would reduce the risk, such as parking in a garage or reducing the time the car is left unattended. When the practical stakes are high, it can be worth spending the time and mental energy to deal with a partial rather than full belief. This follows from the general principle that evidence-gathering effort,

postulated by some philosophers (Williamson 2000: chapters 9 & 10; Fantl & McGrath 2009: chapters 1 & 7).

computational resources and time are allocated where they are needed most, and thus where accuracy is most important (Simon 1956, Payne & Bettman 2004, Bogacz 2007).

This theoretical prediction is confirmed by experiments showing that when the stakes are raised, people collect more evidence before taking a matter to be settled. Judging whether the tendency of the dots on a screen is to move to the left or to the right is a ‘perceptual choice’. Participants are given a fixed amount of time, say two minutes, to complete an unlimited number of such decision tasks. They are given a monetary reward for each correct answer. The stakes can be manipulated by imposing a ‘time out’ penalty for wrong answers—four seconds before the next task starts, reducing the opportunity to earn rewards. When the stakes are increased, people slow their decision-making, increasing their accuracy, thereby improving their rewards (Simen et al. 2009, Bogacz et al. 2010, Balci et al. 2011). Telling them to concentrate on being accurate has a similar effect (Ratcliff & McKoon 2008). In perceptual choice tasks, people face a ‘speed-accuracy tradeoff’ (Wickelgren 1977, Chittka et al. 2009, Heitz 2014), and the balance they strike depends partly on how much it is at stake in a given decision.

In the cognitive and computational neuroscience literature, perceptual choices are typically modelled as the accumulation of evidence for each option to a threshold, possibly from a biased starting point. (Bogacz 2007, Bogacz et al. 2007, Gold & Shadlen 2007, Ratcliff & McKoon 2008, Wagenmakers 2009, Heitz 2014, Standage et al. 2014.) It is plausible that other types of judgement and value-based decision are arrived at by neural mechanisms implementing this type of algorithm. Some of these processes produce a preference or choice of action; others produce categorical states of judgement, not just degrees of belief.⁶ The rate of evidence accumulation is experimentally distinguishable from the threshold. When the parameters in these models are fitted to experimental data, the results entail that people adjust their decision thresholds according to the stakes and

⁶ One might object that ‘perceptual choices’ are not outright judgements, but plans to press one button or the other. But one can judge that the dots are mostly moving to the left; and one can act on the basis of such a judgement, say by pressing a button. That’s how you’d approach a one-off instance of the task without time-pressure or reward. We’d need good reason to think things are different when people participate in studies concerning perceptual choice.

time constraints (amongst other factors), in an optimal way.⁷ A general picture emerges, according to which the stakes in the task at hand affect the amount of evidence a person requires for judgement.

This literature converges with two others reviewed by Nagel (2010 §§1–2). First, there is the ‘Judgement and Decision Making’ (JDM) sub-field of cognitive psychology. Psychologists working in this area concur that increasing the stakes can increase the strength of evidence we will require for judging. They disagree about the mechanism by which the effect is achieved. I don’t need to take a stance on that debate, but I will note that the evidence increasingly favours ‘single-mechanism’ models, such as a shifting evidential threshold for full belief, over models where we shift between different strategies for answering a question (Newell & Lee 2011, Glöckner et al. 2014, and Söllner et al. 2014). That is, the JDM literature converges with the neuroscience reviewed above.

Second, a literature in social psychology posits that in a given situation, a person has a particular degree of ‘need for closure’ on a question (Kruglanski & Webster 1996). ‘Closure’ on a question is coming to a settled answer. The lower one’s need for closure, the more evidence one requires for judgement; and fear of a costly mistake will cause a lower need for closure (Kruglanski & Webster 1996: 264). So the ‘need for closure’ construct predicts that raising the practical stakes will cause people to require more evidence for judgement. Mayseless & Kruglanski (1987) and Roets et al. (2008) confirmed this prediction experimentally. Mayseless & Kruglanski (1987: 167) raised the stakes by offering extra credit to the student participants for getting at least 9/10 correct on their tasks. Roets et al. (2008: 786) raised the stakes by telling participants they would be informed whether their performance placed them in a ‘skilled’ or ‘unskilled’ group.

Let me introduce some terminology. One employs different ‘implicit standards’ for occurrent full belief that p depending on what’s at stake in the task one ‘calibrates to’. One’s ‘implicit standard for belief that p ’ just describes one’s dispositions to form or withhold full belief in p given different evidence. One ‘calibrates to’ a particular task iff one’s representation of the task causally sets one’s implicit standards for full belief; i.e., the

⁷ Standage et al. (2014 §4) argue that stakes also affect how evidence is encoded, and how it is then integrated.

task causes one to regulate the formation of occurrent full beliefs in a certain way. (For how to understand ‘how much is at stake’ in a task, see Gerken (2017: 40–1, 133–5, 264–5), Turri et al. (2016: 214) and Francis et al. (2019: 430–1).)

Don’t assume that at every moment there is a uniform standard that one applies to all one’s occurrent beliefs. I’m just saying that the perceived practical stakes can make a difference to whether the subject forms an occurrent full belief in a given proposition, or a merely partial belief. The standards are not explicitly represented by subjects; they are implicit in the subject’s dispositions to form occurrent full or partial beliefs. People’s standards shift when they calibrate to a particular practical task, but there may well be a default, general purpose standard. Indeed, I will assume there is.

People’s standards for belief in a particular proposition will often shift when they consider different practical questions. Consider the claim that my car won’t get stolen from Pine Street if I park it there every weekday all year. I might affirm this claim when deciding whether it is worth parking on Pine Street and walking an extra three blocks to the office, which is on a dodgier road. Yet I might suspend judgement on the claim when deciding to pay a few dollars for a year’s anti-theft car insurance.⁸

Given the theoretical computational arguments, the empirical work in neuroscience, Judgement and Decision Making (JDM) and social psychology, and the

⁸ DeRose (2009: 270–3) raises a case in which the subject acts on two practical questions separately but at once, using incompatible attitudes. Jane’s employer offers her a year’s free life insurance. She walks over to the benefits office to sign up—after all, you never know what might happen. On the way, Jane answers her cell phone. It’s one of Jane’s friends, asking whether she’s going to look for a different job in the near future. Jane says she won’t—she’ll still be working at the same place this time next year. She carries on walking to the benefits office for life insurance while occurrently judging that next year she will still be working at the same place, and thus assuming she won’t die in the meantime. She is able to keep separate the two settings in which the question arises as to where she’ll be in a year’s time. If Jane is asked why she is walking to the benefits office, she’ll revert to suspending judgement on whether she’ll be alive and working there this time next year. Jane is able to switch back and forth between calibrating to one practical task or the other, and thus the appropriate attitude to whether she’ll be working in the same place next year. This psychological ability is crucial to computationally efficient reasoning.

pre-theoretical appeal of the claim, it is a well-motivated hypothesis that people ‘calibrate’ to practical tasks, affecting their ‘implicit standards’ for full belief in task-relevant propositions. But given this, it is compelling that we will sometimes calibrate to a practical task for someone else, in order to help them. That is, the helper must regulate their full beliefs in a manner appropriate to the task they are helping with. Let’s see how this works in a variant of the bank case.

Suppose Adam is helping Hannah and Sarah gather information relevant to their task, namely paying in their crucial cheque before Monday while minimizing queuing at the bank. Adam gets evidence Hannah and Sarah don’t have: Bert tells him that since the bank was open on a Saturday three months ago, it will be open this Saturday too. Adam should continue investigating whether the bank will be open this Saturday. That is, he shouldn’t form a full belief that the bank will be open. And he doesn’t form an inquiry-terminating full belief, because he calibrates to the practical question for Hannah and Sarah, and thus regulates his full beliefs in an appropriately stringent way. The absence of a full belief also stops Adam from telling Hannah and Sarah that the bank will be open, which would be unhelpful.

Helpful information-sharing is a central aspect of human cooperative behaviour (Tomasello 2009, Sterelny 2012). Pre-verbal children can’t help themselves but to inform others of what they want to know, by pointing at the object someone is looking for (Tomasello 2009: 14–21, Warneken 2015: 2). People have the ability and motivation to helpfully share information. It is a sensible hypothesis that, as part of that ability, we sometimes calibrate to a practical task for someone else.

When will someone calibrate to a practical task for someone else? We don’t do so *whenever* we become aware of a task someone is pursuing. For one thing, doing so would make it hard to gather testimony without shifting our calibration away from the task that is motivating our search for information. For another, we don’t treat everyone we hear about as a friend in need—we are naturally parochial in our affections (Hare 2017: 170). No doubt many factors affect calibration, but I will suggest one cue. I suggest that people are more likely to calibrate to a task when it is ‘psychologically close’ to them. Tasks are psychologically close when they are tasks in the here-and-now for someone socially related; these are the tasks that are ripe for helping with. The next section explains

psychological distance and its typical effects. There is good independent evidence that psychological closeness is a cue to mentally prepare for cooperative action. Thus it is a sensible hypothesis that it performs the same function in the case at hand. That is, the psychological closeness of a practical task cues us to calibrate to it. That's a claim about how people regulate their formation of full beliefs. Building on that claim, §4 formulates my hypothesis about how we ascribe knowledge.

3. Psychological background: psychological distance and its effects.

Cognition is a means for dealing with one's environment—"thinking is for doing" in Susan Fiske's slogan (1992). If we are tackling a practical problem, we construct the mental representations that are relevant and appropriate for dealing with it. The literature reviewed in the last section illustrates this theme: we tailor our full versus partial beliefs to the stakes in the task at hand. For beings like us, the *social* environment is as important as anything. We construct mental representations that help us act appropriately in the particular social environment we find ourselves in (Tetlock 2002).

Does this mean our mental states will be stable over time, or variable? The contemporary answer is: it's a bit of both (Ledgerwood 2014: 436–8). We've already seen that it is useful to vary one's full beliefs depending on how much is at stake in the practical question one is addressing (§2). This section explains why it can be functional to vary our evaluations, by helping us collaborate with others.⁹

Humans evolved to cooperate. All other primates pay a metabolic cost to colour their sclera, disguising the direction of their gaze; but humans have white sclera and almond-shaped eyes that broadcast what we are looking at, setting us up to work together rather than grab food for ourselves (Kobayashi & Kohshima 1997). There's disagreement about the psychological capacities that produce collaborative behaviour, and whether they are the products of genetic or cultural evolution. Tomasello (2009) and Hare (2017) emphasise narrowly-defined capacities and genetic evolution, Heyes (2018) general capacities and cultural evolution, and Sterelny (2012) the interplay between the

⁹ Other arguments that cognition is sensitive to the social context include: Smith & Semin (2004), Schwartz (2007), and Yeh & Barsalou (2006).

aforementioned factors. These are all accounts of why humans are uniquely collaborative among primates.

To be successful, humans must form groups in which they work together. Collaborating helps us achieve a better outcome now, and cements membership of the group, setting us up for better outcomes later. Working together requires common beliefs and common goals—‘coordination’ on the relevant matters. So we should expect people to align their beliefs and evaluations with those of others with whom they might cooperate. And that is what psychologists find.

Let me give you the flavour of Alison Ledgerwood’s (2014) survey of some relevant empirical work. In a series of studies (pp. 439–441), participants were informed of a government policy, such as more vigorous deporting of undocumented immigrants, and were told either that the policy was to come into effect next week, or next year. Participants had a brief conversation with a stranger who expressed an opinion about the policy, and then were asked to evaluate the policy privately. If the policy was to be implemented the following week, then participants shifted their evaluations towards those of the strangers they happened to meet. If the policy was to be implemented next year, participants did not shift their evaluations towards those of random strangers. However, they did shift their evaluations towards those of society at large, or of some salient stable group they identified with.

This is as predicted by need to cooperate. If we are preparing to deal with an event happening soon, then we have to cooperate with whoever is close by, even a random stranger. Cooperation requires a measure of agreement, and so we shift our views towards those of whoever is close by. But if we are considering an event happening in the distant future, then it’s unlikely we’ll have to cooperate with this stranger then; more likely we’ll be cooperating in one of our stable groups.

These findings are part of a more general pattern. The above experiment manipulated participants’ temporal distance from the event they evaluated. But temporal distance is just one dimension of an event’s ‘psychological distance’ from a thinker. (Trope & Liberman (2010), Fiedler et al. (2012), Liberman & Trope (2014), Ledgerwood (2014), the meta-analysis of Soderberg et al. (2015) which integrates over 200

experiments, and Van Boven & Caruso (2015).) The four standard dimensions of an event's psychological distance are:

1. The event's felt temporal distance from the thinker.
2. Its felt spatial distance from the thinker.
3. Hypotheticality: (i) real events feel closer than hypothetical ones; (ii) probable events feel closer than improbable ones.
4. The felt social distance of the people involved. Socially distant people are: not oneself; not similar to oneself; not familiar; not part of one's ingroups (the groups with which one identifies).

One event from last year feels as if it were yesterday; another feels a lifetime ago. A disaster on the other side of the world feels remote, as do the misfortunes of local homeless people.

Psychologically close events tend to evoke stronger emotion and motivation; they call for action now (Peetz, Wilson & Strahan (2009), Van Boven & Caruso (2015: 593 & 597–8), Nook et al (2017).) Psychological distance is automatically computed. It exerts the following tendencies:

- (a) When an event is psychologically close, people evaluate it in a way that's tailored to the task of doing something about the matter now, and hence task-specific and variable.
- (b) When an event is psychologically distant, people evaluate it in a way that's general-purpose, and hence more stable.

These tendencies are functionally beneficial, because psychologically close events are usually the ones that call for action now. Psychologically distant events usually don't call for action now, and our thinking about them might be used in any number of future tasks in a variety of circumstances. So it is typically better to think about distant events in a general-purpose way.

Knowledge ascriptions are evaluations—they are the most common epistemic evaluations.¹⁰ So we should expect a similar pattern of task-sensitivity versus stability in our knowledge ascriptions. That is: when considering a psychologically close scenario, people ascribe knowledge using standards tailored to doing something about the matter now. When considering psychologically distant scenarios, people ascribe knowledge using a general-purpose, stable standard. This pattern should be functional, and partly because it helps us cooperate with appropriate others.¹¹ I will present an account with these features.

Much work on psychological distance investigates ‘Construal Level Theory’. According to CLT, psychological distance affects our evaluations by affecting how we represent the relevant events. Psychologically close events are represented concretely with lots of specific detail, while psychologically distant events are represented more abstractly, and this explains why we evaluate close and distant events differently. CLT can explain many experimental findings, though apparently not all (Van Boven et al. (2010), Wessler & Hansen (2016), Mrkva et al. (2018: 362)). I don’t need to assume that psychological distance affects knowledge ascriptions solely by affecting the level of construal. Psychological closeness functions generally as a cue to tailor one’s thinking to that particular task, even if this effect is then achieved differently in different cases. As I’ll explain at the end of §5, I present my proposal in a manner incompatible with CLT, though there is a CLT-friendly variant. This softens a worry stemming from the ‘replication crisis’ in social psychology (Open Science Collaboration 2015, Benjamin et al. 2018). I rely on the core claim that psychological closeness is a cue to tailor one’s thinking to that particular task, not on predictions peculiar to the Construal Level Theory of the mechanism by which psychological distance affects evaluations. The empirical motivation

¹⁰ Nagel (2010: 408) cites that ‘knows’ is one of the ten most common verbs in English.

Knowledge ascriptions dwarf evaluations of epistemic rationality or justification, even if only half of uses of “knows” are epistemic evaluations (Hansen, Porter & Francis 2019: 12).

¹¹ Shea et al. (2014) argue that the function of communicable metacognitive representations—presumably including knowledge ascriptions—is to improve complex cooperation. The suggestion is compatible with the concept of knowledge having evolved culturally rather than genetically (Shea et al. 2014: 191, Heyes 2018 chapter 7).

for my proposal is not undermined if predictions peculiar to CLT are overturned (e.g. Trautmann 2019). Nevertheless, we should recognize that because my hypothesis is empirically supported, it is also empirically vulnerable. If it turns out that there's no causally relevant factor in the vicinity of psychological distance, then the hypothesis is wrong.

4. The psychological hypothesis.

My hypothesis about how people ascribe knowledge consists of the following claims:

- (1A) Sometimes, people 'calibrate to' a practical task (for themselves or for others), and thus govern their occurrent full beliefs by implicit standards that are appropriate to that task. Other times, people don't calibrate to a task, but employ a general-purpose default standard.
- (1B) The more psychologically distant the contemplated practical situation, the more likely the judge is to stick with the default implicit standard for full belief. The psychologically closer the contemplated practical situation, the more likely the judge is to calibrate to it.
- (2) One ascribes knowledge partly by assessing the subject's evidence using one's own current implicit standards for full belief.
- (3) People typically use and interpret the word "knows" to mean: knows.

Claims (1A) and (1B) summarize §§2–3. They extrapolate naturally from orthodox views in psychology. (1B) is a *ceteris paribus* generalization, compatible with other factors affecting what judges calibrate to (this will be relevant in §8). (3) just means that there's no systematic linguistic funny-business going on, either semantic or pragmatic. Let me say a few words about (2).

Claim (2) ties knowledge ascriptions to what (1A) and (1B) say about people's implicit standards for full belief. Roughly: if one would not form a full belief on the basis of certain evidence, then one will judge that such a belief does not constitute knowledge. On the other hand, if one's implicit standard licenses a full belief, then (other things being

equal) one will judge that such a belief does constitute knowledge.^{12, 13} Of course a belief must meet other conditions to be assessed as knowledge. If one judges that p is false, then one won't judge that the subject knows that p , no matter how strong their evidence that p . Gettier cases must also be accounted for. But those features of knowledge ascriptions are plausibly accounted for separately (as they are according to Powell et al. 2015). So let's set those other features aside for this paper, and concentrate on explaining how and when knowledge ascriptions are sensitive to practical stakes.

I will illustrate the hypothesis by considering how three kinds of people ascribe knowledge in the bank case: Hannah, friends of either Hannah or Bill, and participants in fixed-evidence X-phi surveys (§§5–7 respectively). I will explain why the predicted knowledge-ascriptions are the useful ones for those judges to make. So are the predicted judgements about what others mean by “knows”. The hypothesis sits well with the axiom that thinking is for doing, which applies even to thinking about what someone's words mean.

The hypothesis is well-motivated by the following three considerations. First, the hypothesis plausibly extends widely-held theories in psychology; it is well-motivated by what we know about the mind. This is most true of elements (1A) and (1B). Second, the hypothesis issues in predictions that seem right (§§5–7). In the case of fixed-evidence

¹² Nagel (2012: 186) anticipates claim (2). She suggests that “our own way of thinking” provides the “standard for evaluating the subject's thinking”. There's precedent for the idea that the mechanisms regulating one's beliefs are central to evaluating the beliefs of others. Plausibly, the mechanisms regulating how we act play a role in evaluating the actions of others. According to Miller & Cushman (2013), we sometimes evaluate others morally by simulating acting in the relevant ways ourselves, and seeing how that would make us feel.

¹³ Turri, Buckwalter & Rose (2016) argue empirically that judgements about what someone should do (‘actionability judgements’) have a strong causal impact on our ascriptions of knowledge to them. According to (2), actionability judgements and knowledge ascriptions are instead sensitive to a common cause, namely one's implicit standard for full belief, which is sensitive to the demands of the task one calibrates to. Turri et al. admit that their results can be well accounted for by such a common cause (2016: 220). They measure assertibility judgements rather than actionability judgements generally, which might be a special case (§8).

survey participants, we know the predictions are correct. We currently lack experimental confirmation of the other predictions, a situation that needs to be remedied. Until then, I leave it to the reader's conscience to determine how strongly to trust their sense of what Hannah's friends would think. Third, the knowledge ascriptions the hypothesis predicts are those most useful given the thinkers' practical and social situations (§§5–7). Given the general axiom that thinking is for doing, this lends credibility to the predictions, and hence to the hypothesis. Taken individually, each of these three considerations motivate the hypothesis only modestly. But when three independent witnesses give the same answer, they acquire a new level of credibility.

Two quick remarks about how my discussion relates to the literature. First, existing psychological proposals appear to vindicate classical invariantism. According to the hypotheses of Nagel (2010, but not 2012¹⁴) and Gerken (2017 chapter 12), stakes effects on knowledge ascriptions result from quite general cognitive biases. Stakes effects thus look like mistakes, vindicating classical invariantism. Dinges (2019 chapter 8) also argues that his psychological proposal is hospitable to classical invariantism. By contrast, I start with the axioms that thinking is for doing and action is often cooperative, and find that our knowledge ascriptions shift in functionally beneficial ways. Nagel (2010), Gerken and Dinges' psychological hypotheses suggest an error theory of stakes effects, whereas mine suggests a positive assessment (§9 below). Contrary to the pattern in the existing literature, serious psychological hypotheses need not favour classical invariantism. (I don't want to get into the details of other proposals here, but the following footnote cites potent criticisms of Nagel and Gerken.¹⁵)

¹⁴ Nagel (2012) repudiates her (2010) psychological hypothesis, but still seems to be assuming classical invariantism. Considering possibilities of error makes us deny that more innocent subjects have knowledge. For example: if we consider the counterfactual possibility that the clock had stopped, then we judge that Wanda's glance at the clock is not enough for her to know that it is 4:15pm (Nagel 2012: 174–5). Nagel claims that we must reject these intuitions on the pain of general skepticism (p. 187). But that's not right according to contextualism, nor the bolder relativism I favour (§9).

¹⁵ Nagel (2010) and Gerken (2017 chapter 12) argue that their hypotheses are well-motivated by existing psychological theories, but that's been challenged vigorously. Dimmock (2019) criticises

Second, there is a flowering literature on the social function of knowledge ascriptions, springing from Craig (1990), such as the essays in Henderson & Greco (2015). Some of this work addresses bank cases, typically arguing for contextualism, e.g. Greco (2008), Fricker (2008), Henderson (2011), Hannon (2013), and McKenna (2013). In my view, the social function of knowledge ascriptions favours a relativist-friendly psychological view, not a contextualist-friendly view (§6). That is, it is typically more useful to treat each other as talking about the same thing, namely knowledge. That's just as well, given that the results of experimental philosophy count heavily against contextualism (§7). While the work influenced by Craig often defends a general theory of 'the' function of knowledge ascriptions, my claims are more modest. I focus on particular situations, asking which mental representations would help those agents negotiate their predicaments.

5. How Bill and Hannah ascribe knowledge.

Let's apply the hypothesis just outlined to the bank case (§1). Our hypothesis predicts that Bill will judge that the bank will be open. When questioned by Hannah, Bill is unaware of the particular use his information will be put to. So he employs the default implicit standard for full belief—or maybe one appropriate to most tasks performed at the bank,

Nagel's appeal to 'egocentric bias', and Dinges (2019: 121–5) criticises Gerken's appeal to 'attribute substitution'.

Nagel's (2010) proposal is not designed to explain the X-phi results, and the prospects strike me as dim. Because it appeals to 'egocentric bias', Nagel's theory implies that people with higher 'need for cognition' will show weaker stakes effects (Alexander, Gonnerman & Waterman 2014). Philosophers have higher need for cognition than most people. So Nagel's theory predicts that philosophers will show weaker stakes effects than ordinary people. That's backwards: X-phi finds that ordinary people have stakes insensitive responses to vignettes, while many philosophers report stakes sensitive intuitions (§§7–8). It would be good to solidify this argument by testing whether need for cognition correlates with an individual's stakes sensitivity, like Alexander et al. do for sensitivity to salient possibilities of error. Nagel (2012: 182–8) gives other reasons to doubt her (2010) hypothesis.

for which it won't be terribly important that they be completed before Monday. Either way, Bill's evidence meets this threshold.

Our hypothesis predicts that Hannah will not take Bill's testimony to settle that the bank will be open. She calibrates to a particular practical task, namely paying in her cheque before Monday. It's natural for Hannah to calibrate to that task, partly because it is psychologically close. The task is one for her and Sarah, not someone socially distant. It concerns a choice she will really make, not a merely hypothetical one. And the choice is to be made here and now. So it is in line with (1B) that Hannah calibrates to the task of paying in her cheque on time. There's a lot at stake for Hannah, so she employs a high implicit standard for judging that the bank will be open on Saturday. Thus she does not form a full belief on the basis of Bill's testimony. Bill reports his evidence, and it falls short of Hannah's threshold for full belief.

People's implicit standards governing full belief also govern their knowledge ascriptions (says 2). So Bill judges that he knows the bank will be open: his evidence is enough for a full belief. But Hannah won't form a full belief on the evidence that the bank was open on a Saturday two months ago. Thus she judges that's not enough evidence to know. In particular, Bill doesn't know. According to (3), Bill and Hannah will take each other to mean the same thing by "know", namely knowing. Hannah has no reason to think Bill is lying; she takes him to be expressing something he really believes, namely that he knows. But he doesn't know. So the hypothesis predicts that Hannah interprets Bill as wrongly thinking that he knows. According to Hannah, Bill only thinks he knows because he's being too lax.

The psychological hypothesis seems to give the right predictions about what Bill and Hannah will think. Moreover, that's the only relevant and useful way for Hannah to assess whether Bill knows. Given Hannah's task, Bill's testimony doesn't allow her to settle that the bank will be open. So she should judge that Bill's belief falls short epistemically. Any other evaluation would be a useless distraction, a waste of cognitive resources. Given that 'thinking is for doing', we should expect Hannah to make the relevant evaluation and no other. If Hannah assesses Bill epistemically using the concept KNOWS, then she will judge that Bill doesn't know, as my hypothesis predicts.

If Hannah judges that Bill doesn't know, as my hypothesis predicts, then classical and interest-relative invariantism condemn her. Those views demand that Hannah *not* assess whether Bill knows using the high standard that's relevant to her practical situation. Classical invariantism demands that she assess whether Bill knows using the one true and universal standard. Interest-relative invariantism demands that she use the low standard that's relevant to Bill's practical situation.¹⁶ Is it tenable to concede that Hannah will judge that Bill doesn't know, as is practically useful, but insist that she'll be wrong? I'm saving my invective against such views for §9.

(As I explained at the end of §3, some psychologists subscribe to the Construal-Level Theory of how psychological distance affects evaluations. CLT says that psychological distance affects evaluations by affecting the way the targets are represented (specifically, the degree to which they are represented concretely versus abstractly). The evaluation is performed in the same way, just on differing input representations, and hence with differing results. That does not fit with how I've been talking about how people assess Bill's belief. To see this, suppose Hannah evaluates Bill while calibrated to her high-stakes task, and Lucy evaluates Bill while calibrated to her low-stakes task (Lucy's cheque isn't important). Bill is similarly psychologically distant for both Hannah and Lucy—that's why they don't calibrate to his practical interests. I say that Hannah and Lucy bring different epistemic standards to the task of assessing whether Bill knows, and that's why they come to differing judgements. CLT implies that Hannah and Lucy apply the same epistemic standard to Bill, but give differing answers because they represent him differently. Lucy evaluates a representation of Bill's situation that includes

¹⁶ Could IRI weasel out of this indictment, claiming that talking to Hannah puts Bill in a high-stakes situation without him realizing? Then Bill doesn't know, as per Hannah's useful evaluation. (Stanley 2005: 118–9 considers this manoeuvre.) This weaselling doesn't solve the problem. Suppose that not much is at stake for Amy, and Bill tells her that the bank will be open on Saturday morning—he was there on a Saturday a couple of months back. Hannah eavesdrops. Hannah should still think Bill doesn't know. But surely Hannah's interests don't trump Amy's in defining Bill's practical situation. The deep problem here for IRI is that people should be able to assess whether Bill knows using whatever standard is relevant to *them*. The next section brings this out by considering the useful knowledge ascriptions of both Hannah's friends and Bill's friends.

the fact that not much will happen to Lucy if her cheque is not paid in immediately. That fact gets into the representation Lucy evaluates because Lucy's cheque-paying is psychologically close to her. Hannah evaluates a representation of Bill's situation that includes that her cheque *is* high-stakes for her. Thus the stakes in their respective bank-tasks affect whether they judge that Bill knows the bank will be open. Rather than getting bogged down in whether my version of the proposal is superior to the CLT-friendly version, let's set the question aside for now and keep moving.)¹⁷

6. How Hannah's friends and Bill's friends ascribe knowledge.

It is human nature to help people, especially by sharing information (Tomasello 2009: 6–21). We are especially eager to help our friends. Suppose Hannah and Sarah inform their friend Adam of their predicament. Adam wants to help by establishing—to the standard relevant to his friends' practical task—whether the bank will be open. It is in keeping with (1B) that Adam adopts a stringent implicit standard for full belief that's appropriate to Hannah and Sarah's practical question. His friends are socially close to him, and the choice they face is real and in the near future. Thus their practical question is psychologically close to Adam. Given that he adopts a similar implicit standard for full belief as Hannah and Sarah, Adam will also ascribe knowledge in the same way (says 2). If informed of Bill's testimony, Adam will think that Bill doesn't know that the bank will be open. Like Hannah, Adam will interpret Bill's utterance as indicating that Bill believes he knows (says 3); but Bill only thinks that because he's being too lax.

¹⁷ Even if CLT is not the right account of stakes effects, it may be the right account of other effects on information search and knowledge ascriptions. For example, Halamish & Liberman (2017) find that people underestimate how annoying it is to sample marbles from two urns 100 times, when imagining doing it in a year's time, or someone else doing it. Halamish & Liberman give a CLT explanation: when considering those psychologically distant scenarios, people represent the sampling process abstractly, rather than in concrete detail. That's compatible with my proposal. However, they say CLT supports a general claim: "Psychological distance from a decision situation makes people construct larger data bases from which to draw conclusions." (2017: 111.) If there is such a tendency, it must typically be swamped: surely people don't expend more cognitive effort the less action-relevant a question is for them.

This is the only relevant and useful way for Adam to ascribe knowledge. To help Hannah and Sarah, Adam must adopt the implicit standard for full belief that's appropriate to his friends' situation. He'd be useless to them if he formed a full belief that the bank will be open on the basis of Bill's testimony—he'd take the matter to be settled, rather than needing further investigation he can help with. It wouldn't be helpful to have a laxer standard for full belief than is appropriate to Hannah and Sarah's task, for then Adam would close investigation too soon. Nor would it be helpful to have a more stringent standard, for then he would continue investigation too long. To be useful, Adam's knowledge-ascriptions must be driven by the task-appropriate implicit standard for full belief. And that's what the psychological hypothesis predicts Adam does. So again, the hypothesis comports with the axiom that thinking is for doing.¹⁸

Now consider Bill's friend Katie. They are going to hang out together on Friday evening. But should Bill first spend an hour of his afternoon in line at the bank, paying in his cheque? There's no need, Katie will think; Bill knows the bank will be open tomorrow. Some stranger (Hannah) said that Bill doesn't know?—Well then she's being excessively cautious. The psychological hypothesis predicts Katie's responses in the same way it predicts Adam's. Again, this is the useful way for Katie to ascribe knowledge, given her social situation, namely her focus on Bill and how he should start his Friday night.

In sum: Adam sets himself to cooperate with Hannah and Sarah, while Katie sets herself to cooperate with Bill. Thus they adopt implicit standards for full belief relevant to their respective friends, and ascribe knowledge accordingly. That's the useful way for

¹⁸ Many philosophers are attracted to principles like RKP:

(RKP) S can take p to be available as a reason for acting iff S knows that p .

(See Hawthorne & Stanley 2008, Jackson 2012.) An anonymous reviewer worries that my hypothesis predicts that people will accept counter-examples to RKP. That is: people would regularly make judgements of the form, "S doesn't know that p but S should still act on p ", and, "S knows that p , but S shouldn't act on p ", which they don't. My hypothesis has no such consequence. Verdicts about whether S can take p to be a reason shift in lock-step with verdicts about whether S knows that p . Both are evaluated using an epistemic standard set by the task the judge calibrates to. (Considering what someone ought to do may often cause us to calibrate to their task—see the end of §8.)

them to think and talk about knowing, which makes it psychologically plausible that they do so. My hypothesis again makes the right predictions, both from the perspective that ‘thinking is for doing’, and pre-theoretically. Conversely, it is psychologically implausible that Adam would ascribe knowledge to Bill in the way deemed correct by classical or interest-relative invariantism, because such a judgement would be irrelevant to Adam’s situation.

If Adam and Katie judge as my hypothesis predicts, then classical and interest-relative invariantism condemn at least one of them. It is useful for Adam to assess whether Bill knows by one standard, and for Katie to assess it by another. At least one of those assessments is prohibited, according to those philosophical views. Those views hold that there can only be one correct standard by which to assess whether Bill knows, which Adam and Katie must both employ.

Contextualism permits Adam and Katie to say whether Bill “knows” by the standard that is relevant to them. However, it demands that they evaluate Bill’s utterance of “I know” by a different standard. Contextualism demands that everyone evaluates the truth-value of Bill’s utterance of “I know” using the epistemic standard determined by Bill’s conversational context. Assessing Bill’s utterance in that way will typically be irrelevant to the task at hand, an unhelpful distraction and a waste of cognitive resources. Adam calibrates to helping Hannah and Sarah, and so the *only* relevant evaluation of Bill’s belief is that it falls short. The only evaluation of Bill’s belief that’s relevant for Katie is that it is good enough. For both Adam and Katie, it would be an unhelpful distraction to also figure out what standard was appropriate to Bill’s conversational context and assess Bill’s belief against it, to judge whether Bill’s utterance was true. So contextualism condemns the useful way of evaluating the truth-value of utterances ascribing knowledge. It demands that people evaluate the truth-value of such utterances in a task-irrelevant and useless way. Given that thinking is for doing, it is psychologically implausible that they do so.

One might reply that hearers must evaluate Bill’s claim to “know” by the lax epistemic standard he employs, so as to draw the appropriate conclusions from his assertion. Suppose Bill hadn’t told Hannah what evidence he possesses that the bank will be open. Then for Hannah to avoid relying on Bill’s testimony, she would have to

understand that Bill is using a lower epistemic standard. So, goes the worry, it isn't an unnecessary distraction to assess Bill's utterance using his lower standard, as I have alleged. It is useful to assess knowledge-ascribing utterances in the way contextualism demands after all. But this reply is mistaken. Sensitivity to Bill's use of a laxer epistemic standard does not require employing it oneself when evaluating his utterance. Hannah should think: Bill thinks he knows the bank will be open, but he's being sloppy, too lax. She thereby grasps that Bill takes a lower strength of evidence to be sufficient for knowing; but she still assesses Bill's self-ascription of knowledge using her stringent standard.

Admittedly, Hannah or Adam could search for a more sympathetic understanding of Bill, and thus think that he accurately evaluated his belief using a lower epistemic standard. But this requires calibrating to Bill's practical situation. My point concerns how one thinks about Bill when calibrated to helping Hannah and Sarah.

My argument against contextualism fits a wider trend. Contextualist-style views are wrong for evaluative utterances more generally, not just epistemic evaluations of whether someone "knows". Ledgerwood (2014) reports shifts in participants' evaluation of government policies, such as deporting more undocumented immigrants (§3 above). We should not explain those shifts by saying that participants changed what they meant by a "good" policy. Again, we should not build speakers' moral standards into the contents of their moral judgements and utterances. For example: Aristotle thought that slavery is morally permissible, but he was wrong. We don't ascribe true contents to his 'moral' thought and talk. Why is this the useful way to interpret and assess evaluative utterances? What's important is our evaluation of slavery, knowing Aristotle's evaluation, and that those evaluations clash. So it's useful to think that he was wrong about slavery. There's nothing useful about thinking his 'moral' judgements and utterances were right, by evaluating them in a way alienated from our moral standards. This is the same complaint I raised against contextualism about knowledge ascriptions. The point about knowledge ascriptions gains force because it is an instance of an attractive general claim about how we interpret evaluative utterances.

The objection to contextualism also applies to 'WAM-ing invariantism'. On this view, "knows" always semantically expresses knowing, which requires a moderate

strength of evidence. But in many conversations, it would be irrelevant to talk about that epistemic standard—only a higher one will do. Speakers can then use “knows” to mean the more relevant epistemic standard, trusting that hearers will interpret them charitably. The communicated meaning is conversationally implicated, in the normal way explained by Grice’s conversational maxims (Grice 1989: chapter 2), it is alleged. People’s pattern of assertions using ‘knows’ is explained by a ‘Warranted Assertibility Manoeuvre’, hence the label ‘WAM-ing invariantism’. Proponents include Rysiew (2007) and Brown (2006). DeRose (2009: chapter 3), Blome-Tillmann (2013) and Dinges (2019: 67–74) criticise the view on linguistic grounds. The following objection is of another kind.

WAM-ing invariantism demands that everyone assesses Bill *belief* about his epistemic position using the standard he had in mind. Similarly, what he *meant* when he said “I know”—the claim he conversationally implicated—must be assessed by the epistemic standard he had in mind. (It’s not terribly interesting that we can also assess a knowledge ascription that’s what Bill literally said but didn’t mean.) And so, like contextualism, WAM-ing invariantism condemns Hannah if she judges that Bill thinks he knows, but he’s wrong. Rather, the view requires her to judge that Bill thinks he meets at least low epistemic standards, and he’s right. It requires Hannah to evaluate Bill’s epistemic self-assessment (what he meant and what he believes) in a way that’s irrelevant to her task. It’s not psychologically plausible that she does so, given that thinking is for doing.

7. How participants in fixed-evidence X-phi surveys ascribe knowledge.

Let’s consider another social situation in which someone might evaluate whether Bill knows—a strange social situation, but one that has been the focus of recent investigation. In many X-phi studies, professors ask their students in large introductory classes to answer questionnaires about fictional cases. Other studies are conducted online, say using Amazon Mechanical Turk, paying people to answer the survey. In fixed-evidence questionnaires, participants read a story, and rate on a Likert scale the degree to which they agree the main character knows. Some participants read a story about someone for whom not much is at stake; others read a story about someone who has the same evidence, but for whom a lot is at stake. (We’ll consider evidence-seeking probes in §8.)

These surveys manipulated what's at stake for the fictional character in the vignette, holding fixed how strong their evidence is (or so I will assume¹⁹). They found that the stakes had little or no effect on whether participants judged that the character knows. "Little or no effect" is a disjunctive claim. Some studies found evidence of a surprisingly small effect, namely Sripada & Stanley (2012)—which Francis et al. (2019: 439) failed to replicate—and Pinillos & Simpson (2014 study 2). I discuss what to say if there really is a small effect in §8. Most studies found no evidence of an effect at all, including Buckwalter & Schaffer (2015, and the studies they review in §1); Turri (2017); Francis, Beaman & Hansen (2019); and Rose et al. (2019), which is a large cross-cultural study.

This is bad news for Interest-Relative Invariantism. The findings undermine arguments for IRI that appeal to intuitions about bank-style cases. Moreover, the findings imply that IRI is counter-intuitive. According to IRI, a big difference in what's at stake for the character has a decisive effect on whether the character knows. Participants do not ascribe knowledge as IRI demands (whether the stakes effect is non-existent or merely very small).

There's bad news for contextualism too—specifically, views according to which what's at stake for the speaker affects the content of "knows" in their mouths. Some studies had the character in the vignette utter, "I know". For example, Rose et al. (2019) used two variants of a 'bank'-style case, one where there's a lot at stake for the character paying in their cheque, and one where there's not, and in both variants the character asserts, "I know the bank will be open on Saturday morning." What was at stake for the character had little or no effect on whether participants judged that their claim to "know" was true. So participants did not assess the characters' utterances of "I know" as

¹⁹ I will assume that participants reading the high- and low-stakes versions of the vignettes attribute the same evidence to the characters. Nagel (2010: 429 n. 6, and p.c.) and Pinillos (2011 §4.3) challenge this assumption. They suggest that when participants read about a character in a high-stakes situation who nevertheless takes themselves to know, participants assume the character must have very strong evidence. If that's right, then fixed-evidence vignette surveys would not show up any stakes-sensitivity in participants' knowledge ascriptions.

contextualism demands (specifically, versions where the stakes for the speaker affect the content). This is a serious blow to the view.

My psychological hypothesis predicts these findings. Participants in X-phi surveys are in a strange social situation, answering questions about fictional characters in made-up scenarios. As I will explain, my hypothesis predicts that—*ceteris paribus*—participants will *not* calibrate to the fictional characters' practical questions. The surveys ask about events that are psychologically distant to participants, and so participants will ascribe knowledge by a general-purpose, default standard. Participants given the low-stakes story and the high-stakes story both employ the default standard, and thus ascribe knowledge is the same way. And that's what experimental philosophers found.

Let's run through the dimensions of psychological distance to confirm that the surveys ask participants about very distant events. Firstly, the stories are fictional, not real, and involve fictional characters. Thus the events are distant along the dimension of hypotheticality. That's enough to make the events psychologically distant; but let's continue to ram the point home. The events are even more distant if the story is improbable, as they tend to be, especially in the high-stakes variant. Driving a truck across a rickety wooden bridge over a yawning canyon (Feltz & Zarpentine 2010: 706), losing the ability to stay in college if there's a single typo in their next essay (Pinillos 2012)—these scenarios will seem improbable to most participants, and are thus even more psychologically distant.

Secondly, the characters are socially distant. They are not the participants themselves, nor are they friends or family members. The characters are not familiar to participants (unlike Sherlock Holmes). Depending on the story, the characters might be dissimilar to the participants, and be putative outgroup members. For example, someone depositing a physical cheque is dissimilar from the average undergraduate participant. Kids these days upload photos of their cheques from their phones. Participants with a bias against homosexuals will be psychologically distant from a lesbian couple, such as Hannah and Sarah in Stanley's version of the bank case. Pinillos (2012) uses a vignette involving a student who turns in a paper containing no typos. Such a student is dissimilar to the average undergraduate survey participant, and sounds like an outgroup member.

Finally, there's nothing in the story to make participants think of the events as happening now or in the immediate future. If anything, not specifying when the events take place makes them feel even more remote and abstract than saying they take place next year. The spatial location is also unspecified, with similar consequences.

In sum, the events described in the surveys are psychologically very distant to participants. Thus the hypothesis predicts that participants will ascribe knowledge to the characters by a default standard for knowing. And the X-phi results show that's what people do (when presented with a fixed-evidence probe).

Fixed-evidence probe surveys find no effect of practical stakes on how people ascribe knowledge (let's suppose). It is tempting to extrapolate, and conclude that stakes don't ever affect people's knowledge ascriptions. That would be a mistake. My psychological hypothesis explains the lack of a stakes effect in the strange social situation studied by experimental philosophers, just as well as the hypothesis that there are never stakes effects, and my hypothesis entails that stakes often do have an effect on how people ascribe knowledge.

Zoom out for a moment. Standard X-phi surveys put participants in a very unusual social situation, and ask them about scenarios of no personal relevance. To extrapolate from those results to a claim about how people judge in the wild, we need a substantive further premise. We'd need to assume that people's practical and social situations do not affect how they make the relevant judgements. As we saw in §§2–3, that's not generally true. Unless we are given a good argument that there are no situational effects in the case at hand, we cannot take vignette surveys to show how people think in the wild. In the terminology of research design, vignette surveys are of doubtful 'external validity': they do not support desirably general conclusions. This is a deep challenge for X-phi as currently practised.²⁰

Hansen (2020 §§4–6) raises this concern about X-phi, warning of three situational effects on judgement. He concludes that we can only draw general conclusions about cognition from ecologically valid experiments. The problem is taken seriously across

²⁰ Traditionalists should not gloat, for a related challenge arises for philosophers who rely on their own intuitions. Why rely on the intuitions one has in the seminar room, if one would give different verdicts in other, more engaged practical and social contexts?

psychology. Smith & Semin (2004: 61–2 & 67) warn of it. For example, choosing hypothetically differs from choosing with real-world consequences. What people think they'd do in hypothetical moral dilemmas doesn't match what they'd choose in real life (FeldmanHall et al. 2012, Francis et al. 2016, Bostyn et al. 2018), and the same goes for choosing more generally (Camerer & Mobbs 2017). The field acknowledges that we can't extrapolate real choices from hypothetical choices. Schilbach et al. (2013) review how neuroscientists revised their experiments into social cognition when they realized that social context has a massive influence. Interacting with someone uses very different brain processes than merely thinking about them at some remove. To really study social cognition, neuroscientists had to make their experiments approximate real life interactions more closely, not just having participants observe videoed people passively. Experimental philosophers will need to change their methods too, at least sometimes, conducting experiments that approximate real-life situations. Standard vignette surveys are quick and cheap to conduct, but they don't always tell us what we want to know. X-phi has an exciting future, using the kinds of experimental designs found in the best work in psychology.

Zoom back in on to knowledge ascriptions. Can we squeeze more juice out of vignette surveys? We could make participants' mums the central characters in the vignettes. If we find a stakes effect on knowledge ascriptions, that suggests that people are more likely to calibrate to socially closer people. We could manipulate the temporal distance of the scenario, we could ask participants how realistic the scenario was, and see if those factors affect judgements. Such experiments might confirm that psychological distance guides calibration, but they can't significantly disconfirm the hypothesis. Plausibly, fictional scenarios are automatically near maximally psychologically distant, so that varying the other dimensions of distance won't have a significant effect.

The make-or-break question is whether people calibrate their knowledge ascriptions variably in real life situations. If people don't ascribe knowledge as predicted in §§5–6, then my hypothesis is refuted. I can't yet see how to test these predictions rigorously: I'm not trained in experimental design, I haven't spent years honing that skill as psychologists do. I've made a theoretical advance, explaining the current data in a way

that’s motivated by what we know about the mind. Now it’s the job of experimentalists to pick up the baton. I’m excited to see the results!

8. Two additional phenomena to be explained.

I designed a hypothesis to account for the findings of fixed-evidence X-phi surveys, while predicting usefully stakes-sensitive judgements in the wild. Two extra phenomena complicate the data to be explained.

First, many philosophers have intuitions about bank cases that are stakes sensitive, unlike typical survey participants. For example, they are often tempted, after reading about Hannah and Sarah’s predicament, to judge that Bill doesn’t know the bank will be open. That’s why there’s a philosophical debate about such examples. Why do philosophers’ intuitions differ from those of X-phi survey participants?

Second, X-phi surveys using ‘evidence-seeking probes’ typically found a stakes effect. These studies ask participants questions like, “What is the minimum number of times the climber needs to inspect the rope before she knows that it is tied securely?” In a low stakes version of the vignette, the climber will only be hanging 1 metre off the ground on the rope; in a high stakes version she will be hanging 100 metres off the ground. These surveys typically find a stakes effect. For example, participants required an average of two rope inspections for knowledge when the stakes were low, and three when the stakes were high. See Pinillos (2012), Pinillos & Simpson (2014), and Francis, Beaman & Hansen (2019), while Turri (2017) did not find an effect. (Buckwalter & Schaffer (2015 §2) try to explain away the typical results, but I am not convinced.²¹) What accounts for this difference from fixed-evidence probes?

²¹ Buckwalter & Schaffer (2015: 211–215) argue that Pinillos’ evidence-seeking study found a stakes effect on something, just not on knowledge. They asked participants high- and low-stakes variants of the following questions about knowing, guessing, or hoping:

How many times do you think Peter has to proofread his paper before he
[knows/guesses/hopes] that there are no typos? ____ times.

They found the same patterns of response for “knows”, “guesses”, and “hopes”. Surely there is a common cause to the common pattern: people interpreted the three stimuli as asking the same question. Buckwalter & Schaffer say that participants answer how many times Peter *has to*

The hypothesis presented in §4 does not explain these two phenomena. This does not refute the hypothesis. The hypothesis of §4 is a *ceteris paribus* generalization, which specifies what typically determines whether someone will calibrate to a particular practical task, yet allows that other factors sometimes trump this default. Cognition is a complex business, and many factors are potentially relevant to a given judgement. Even “people prefer to avoid pain” is just a *ceteris paribus* generalization; it would be daft to hold the hypothesis of §4 to a higher standard. If our two extra phenomena are plausibly cases where atypical factors impinge, the hypothesis of §4 is untroubled.

Let’s start with philosophers’ intuitions. There isn’t much empirical work about what’s unusual about philosophers, but there is a literature on why people’s reactions to fictional narratives differ. Sometimes people become imaginatively immersed in a story, reacting emotionally to things they know are not real. Green & Brock (2000) argue that there are three dimensions of ‘transportation’ by a narrative: high levels of mental imagery, emotional involvement, and cognitive engagement. Green & Donahue (2009) and Van Laer et al. (2014) review the recent literature on transportation. The most relevant study is Johnson et al. (2013). They found that constructing vivid mental imagery causes people to be more emotionally involved in a story they read.

Some participants were instructed to generate vivid imagery from a range of sensory modalities while reading the story—to imagine seeing, hearing, smelling, and touching. Others were told to concentrate on the meanings of the words and sentences used—the ‘verbal-semantic’ condition. The rest were assigned to the control condition, told to read the story as they would for leisure. Participants instructed to construct vivid imagery reported doing so, reported higher levels of empathy for the characters in the story, and were much more likely to help pick up the pens the experimenter ‘accidentally’

proofread his paper, ignoring the subsequent words about knowing, hoping, or guessing. I say that in all three cases, participants answer the question about knowing. Participants will re-interpret the questions about hoping and guessing, because those questions are irrelevant and daft. First, what Peter hopes or guesses is of no import in the scenario. Second, Peter does not need to proofread his paper *at all* before he hopes or guesses it has no typos. The only sensible question in the vicinity is the one about knowing (or the equivalent one as to whether to act on the premise that there are no typos—see note 19). (Compare Dinges 2019: 140–2.)

dropped on the floor. They were twice as likely to help pick up the pens as people in the verbal-semantic condition, and three times more likely than leisure-readers. This shows that vivid imagery has wide effects, not just on how people report feeling about a story.

These findings comport with the general theory that mental imagery functions like a weak form of perception (Holmes & Mathews 2010, Pearson et al. 2015 esp. 598). Imagery uses the same brain mechanisms as genuine perception, and has similar effects. Imagery produces stronger emotions than verbal-semantic thinking, and makes the events represented feel more real and more likely. So we should expect that generating imagery while reading fiction results in emotions more similar to perceiving the events described. We can put the prediction in terms of psychological distance. Perceived events are happening here and now; generating mental imagery is like weakly perceiving the events; so generating imagery causes the imagined events to feel psychologically closer.

Let's apply these findings to people's reactions to fixed-evidence vignettes. People who focus on the meanings of the words and sentences used (the 'verbal-semantic condition') will not be transported by the story, and will have a low level of empathy for the characters. People who generate rich mental imagery of the events described will be transported, treating the events as psychologically closer, and will have more empathy for the characters, entering into their predicaments. Applying the hypothesis of §4: people who generate rich mental imagery are more likely to calibrate to the practical task for the character in the vignette, compared to people who consider the story in a verbal-semantic way. This claim is empirically motivated. To explain why philosophers have stakes-sensitive intuitions, unlike ordinary participants in X-phi studies, we need to add the hypothesis that philosophers generate rich mental imagery when performing thought experiments, whereas ordinary participants consider scenarios in a verbal-semantic way. No empirical work speaks directly to whether philosophers differ in this way. But there are empirically motivated reasons why this might be so. Let's start with some differences between the situations of philosophers and participants, before considering patterns of individual differences.

Being transported into a story disconnects one from immediate reality (Green & Brock 2000: 702), so we'd expect that people will only do it if they lack immediate practical goals, and feel relaxed in their environment. Philosophers in their armchairs fit

these criteria, but participants in X-phi studies do not. Survey participants have occurrent, mentally activated goals involving their immediate environments. People answer questions on Amazon Mechanical Turk to earn money. Student participants taking a survey from a scary professor before a class on an unfamiliar subject also have immediate practical goals. They are trying to complete the questionnaire in the time allotted, and they have social goals: not seeming stupid to the professor, and not seeming weird to their classmates. So participants are unlikely to be transported into the narrative, constructing vivid mental imagery and empathizing with the characters. This fits the observation that people are more likely to be transported into a story if they are reading it because they want to, rather than for a college assignment or to be paid (Green & Donahue 2009: 247). Moreover, having their own goals to calibrate to makes participants inherently less likely to calibrate to the task facing the character in the vignette. By contrast, philosophers thinking about their favourite topics in their armchairs, unharassed by immediate real-world goals, are well-situated to be transported by their thought experiments.

To the situational differences between philosophers and participants, we might add patterns of individual differences. I'd guess that philosophers' training helps them to be transported by thought experiments, but their other typical traits do not. Philosophers are trained to conduct thought experiments by imagining the scenario realistically. In the absence of such training, non-philosophers faced with a written vignette will stick to the given verbal-semantic representation of the scenario. Plausibly so. Relatedly, philosophers' familiarity with the examples helps them to imagine the vignette scenarios fluently and easily, which will make the scenarios feel psychologically closer (Mrkva, Travers & Van Boven 2018).

By contrast, I doubt that philosophers' typical character traits are responsible for their transportation into thought experiments. High 'need for cognition' does not correlate with transportation (Green & Brock 2000: 704, 711, 713). High 'need for affect' does correlate with transportation (Appel & Richter 2010), but I doubt that philosophers seek out emotions more than the general population—rather the opposite. I'd be surprised if philosophers are cross-task more vivid imaginers than most people (Faw 2009).

In sum: philosophers are trained to conduct thought experiments by constructing vivid and detailed mental imagery, and face no situational barriers to doing so. Survey participants are not so trained, are given a verbal-semantic representation of the scenario, and are in a situation inappropriate for constructing mental imagery and being transported into the vignette. Thus philosophers are more likely to construct vivid mental imagery, and so be transported into the vignette, empathizing with the characters, and calibrating to the task they face. So philosophers are more likely to ascribe knowledge in a way that's sensitive to what's at stake for the characters; survey participants typically won't. This explanation is speculative but empirically-motivated.²²

Suppose that, contrary to the trend, fixed-evidence surveys end up finding a small stakes effect. The above account of philosophers' intuitions makes this easy to explain: those surveys manage to transport a few participants. For testing the extended hypothesis, it is unimportant whether fixed-evidence surveys find no stakes effect or a small one.

Let's move on to a tentative hypothesis about evidence-seeking probes. Plausibly, judging what someone else ought to do typically causes people to calibrate to the practical question under consideration. If one addresses what someone else ought to do in order to offer them advice, then one should calibrate to their task. The same goes for predicting what they will do, or practising how to choose if faced by that sort of decision oneself. By contrast, it is easy to see why it can be useful to assess whether someone else knows without calibrating to their task: e.g., we are interested in whether testimony meets the evidential standard relevant to our task, not theirs. In this respect, assessing whether

²² De Smedt & De Cruz (2015) argue that philosophical thought experiments typically don't transport philosophers to the degree that reading science fiction does: thought experiment rarely cause strong emotions (though see FeldmanHall et al. 2012). I've claimed that philosophers—unlike survey participants—are transported enough by certain thought experiments to have stakes-sensitive intuitions about them. These claims are compatible, because it is easier to calibrate to a stranger's task than to care about them. People usually give stakes-sensitive advice to a stranger asking for directions, even without a strong emotional reaction to their predicament.

someone ought to assert is like assessing whether they know, so it should not affect calibration like other assessments of how someone ought to act.²³

Fixed-evidence probes ask participants to evaluate a belief the character already has. But evidence-seeking probes ask what the character ought to do: the climber has to decide how many times to check the rope. Participants evaluate possible courses of action for whether the climber will then meet the goal of knowing that the rope is secure. The practical nature of the question causes participants to approach it in the advice-giving mode, and so to calibrate to the climber's practical situation, giving stakes-sensitive answers. At least, that's one hypothesis. As long as evidence-seeking probes plausibly activate some atypical factor or other, the hypothesis of §4 is not refuted.

9. Metaphysics that endorses our useful judgements.

Suppose that the psychological hypothesis is correct. Suppose, moreover, that the knowledge-ascribing judgements thus described are the most useful mental representations for people to construct. Shall we stick to one of the standard philosophical views, and demand people stop thinking and talking in the useful ways? Uncynical philosophers of such a persuasion would accost friends and write public philosophy blog posts, declaring, "You can't ascribe knowledge like that! Admittedly, it is the most practical way to think and talk, but I can't make sense of it within my approach to metaphysics and philosophy of language!" Such proselytizing would be met with dismissive laughter, and rightly so. Rather than fitting our knowledge-ascriptions to some Procrustean bed, we should look for an approach to metaphysics or semantics that legitimizes them. Our philosophy should endorse the useful ways to think and talk.

Some readers have objected that while generally reliable heuristics are the most useful way to think, it does not follow that their outputs are all correct. This objection misunderstands my argument. I've been assessing how well individual mental representations help thinkers deal with their circumstances. In the cases we've examined,

²³ Turri, Buckwalter & Rose (2016) use judgements about what someone should assert (e.g. write in their report) as their examples of judgements about what someone should do ('actionability judgements'). I've given a functional reason to think assertibility judgements and actionability judgements have different effects on calibration, and thus on knowledge ascription.

the knowledge-ascriptions my account predicts are also the only useful epistemic evaluations for those people to make. I shall continue to make those judgements. I endorse them; I say they are quite proper. Alternatively, one could posit a realm of objective facts about knowledge, and find that many useful knowledge-ascriptions fall short by this metaphysical standard. That theoretical choice is optional, unmotivated, and shows Olympian disdain for human living. I am against it. In the case of knowledge ascriptions, “thinking is for doing” conflicts with “thinking is for representing objective facts”; I conclude that the latter is not generally true.

One might think that truth-relativism is flexible enough to legitimize all the useful ways we attribute knowledge. On this approach, all uses of the word “knows” have the same content. A particular knowledge-ascription can be true relative to one judge but false relative to another. An utterance or judgement is legitimate (in the crucial sense) iff it is true relative to the judge who makes it.²⁴ I am not a truth-relativist. My preferred framework is realist about the metaphysically fundamental facts, but anti-realist about everything else. It is designed to capture the relativity and vagueness in some matters, and the objectivity of others (Jackson 2016, 2019, and a book manuscript in progress). This is not the place to adjudicate between truth-relativism and alternative frameworks. My point is that legitimizing our useful knowledge ascriptions directs us to these bold metaphysical options. My psychological hypothesis motivates a genuine relativism about knowledge.

Acknowledgements

I thank the following for helpful conversations and feedback on earlier drafts: Anne Bosse, Andrew Cortens, Stephen Crowley, Chris Frith, Nat Hansen, Richard Holton, Brian Kierland, Rae Langton, Edouard Machery, Benjamin Marschall, Lucy McDonald, Jennifer Nagel, Ángel Pinillos, Baron Reed, David Rose, William Skylark, Jonathan Weisberg, and two anonymous referees for this journal. I would also like to thank the

²⁴ Jackson (2016: 972–3) argues that the point of contemporary truth-relativism is to ground claims about when assertions and retractions are ‘correct’ or ‘acceptable’ in such a sense. One example is MacFarlane (2014: 103–111 & chapter 8), who gives a truth-relativist account of knowledge ascriptions.

organizers and audiences at the Buffalo Experimental Philosophy Conference (2017), the Joint Session of the Aristotelian Society and the Mind Association (2017), the BARSEA epistemology workshop (2017), and the APA Central Division meeting (2016).

References

- Alexander, Joshua, Chad Gonnerman, & John Waterman. 2014. Salience and epistemic egocentrism: An empirical study. In James Beebe (ed.), *Advances in Experimental Epistemology*. London: Bloomsbury.
- Appel, Markus, & Tobias Richter. 2010. Transportation and need for affect in narrative persuasion: A mediated moderation model. *Media Psychology* 13(2): 101–135.
- Bartlett, Gary. 2018. Occurrent states. *Canadian Journal of Philosophy*, 48(1): 1–17.
- Balci, Fuat, Patrick Simen, Ritwik Niyogi, Andrew Saxe, Jessica A. Hughes, Philip Holmes, & Jonathan D. Cohen. 2011. Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Attention, Perception, & Psychophysics*, 73(2): 640–657.
- Benjamin, Daniel J., James O. Berger, Magnus Johannesson, et al. 2018. Redefine statistical significance. *Nature Human Behavior* 2: 6–10.
- Blome-Tillmann, Michael. 2013. Knowledge and implicatures. *Synthese* 190: 4293–4319.
- . 2014. *Knowledge and Presuppositions*. Oxford University Press.
- Bogacz, Rafal. 2007. Optimal decision-making theories: linking neurobiology with behaviour. *Trends in cognitive sciences*, 11(3): 118–125.
- Bogacz, Rafal, Marius Usher, Jiaxiang Zhang & James L. McClelland. 2007. Extending a Biologically Inspired Model of Choice: Multi-Alternatives, Nonlinearity and Value-Based Multidimensional Choice. *Philosophical Transactions of the Royal Society B*, 362: 1655–70.
- Bogacz, Rafal, Peter Hu, Philip Holmes, & Jonathan D. Cohen. 2010. Do humans produce the speed–accuracy trade-off that maximizes reward rate? *The Quarterly Journal of Experimental Psychology* 63(5): 863–891.
- Bostyn, Dries H., Sybren Sevenhant, & Arne Roets. 2018. Of mice, men, and trolleys: Hypothetical judgment versus real-life behavior in trolley-style moral dilemmas. *Psychological science* 29(7): 1084–1093.
- Brown, Jessica. 2006. Contextualism and warranted assertibility manoeuvres. *Philosophical Studies* 130: 407–435.
- Brown, Jessica, & Mikkel Gerken (eds). 2012. *Knowledge Ascriptions*. Oxford University Press.
- Buckwalter, Wesley, & Jonathan Schaffer. 2015. Knowledge, stakes, and mistakes. *Noûs* 49(2): 201–234.
- Camerer, Colin, & Dean Mobbs. 2017. Differences in behavior and brain activity during hypothetical and real choices. *Trends in Cognitive Sciences* 21(1): 46–56.
- Chittka, Lars, Peter Skorupski, & Nigel E. Raine. 2009. Speed–accuracy tradeoffs in animal decision making. *Trends in Ecology & Evolution* 24(7): 400–407.
- Cohen, Stewart. 1999. Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13: 57–89.
- Craig, Edward. 1990. *Knowledge and the State of Nature*. Clarendon Press.

- DeRose, Keith. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52(4): 913–929.
- 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1*. Oxford University Press.
- De Smedt, Johan, & Helen De Cruz. 2015. The epistemic value of speculative fiction. *Midwest Studies in Philosophy* 39: 58–77.
- Dimmock, Paul. 2019. Knowledge, belief, and egocentric bias. *Synthese* 196 (8): 3409–3432.
- Dinges, Alexander. 2019. *Knowledge, Stakes and Error: A Psychological Account*. Frankfurt a. M.: Klostermann.
- Easwaran, Kenny. 2011a. Bayesianism I: Introduction and Arguments in Favor. *Philosophy Compass*, 6: 312–320.
- 2011b. Bayesianism II: Applications and Criticisms. *Philosophy Compass*, 6: 321–332.
- Evans, Jonathan St. B. T., & Keith E. Stanovich. 2013. Dual-Process Theories of Higher Cognition: Advancing the Debate. *Perspectives on Psychological Science*, 8(3): 223–241.
- Fantl, Jeremy, & Matthew McGrath. 2009. *Knowledge in an Uncertain World*. Oxford University Press.
- 2012. Arguing for shifty epistemology. In Brown & Gerken (eds).
- Faw, Bill. 2009. Conflicting intuitions may be based on differing abilities: Evidence from mental imaging research. *Journal of Consciousness Studies* 16(4): 45–68.
- FeldmanHall, Oriël, Dean Mobbs, Davy Evans, Lucy Hiscox, Lauren Navrady, & Tim Dalgleish. 2012. What we say and what we do: The relationship between real and hypothetical moral choices. *Cognition* 123(3): 434–441.
- Feltz, Adam, & Chris Zarpentine. 2010. Do you know more when it matters less? *Philosophical Psychology* 23(5): 683–706.
- Fiedler, Klaus, Janis Jung, Michaela Wänke, & Theodore Alexopoulos. 2012. On the relations between distinct aspects of psychological distance: An ecological basis of construal-level theory. *Journal of Experimental Social Psychology* 48: 1014–1021
- Fine, Kit. 2001. The Question of Realism. *Philosophers' Imprint* 1/1.
- Fiske, Susan T. 1992. Thinking is for doing: Portraits of social cognition from daguerreotype to laserphoto. *Journal of Personality and Social Psychology*, 63(6): 877–889.
- Francis, Kathryn B., Philip Beaman, & Nat Hansen. 2019. Stakes, Scales, and Skepticism. *Ergo* 6(16).
- Francis, Kathryn B., Charles Howard, Ian S. Howard, Michaela Gummerum, Giorgio Ganis, Grace Anderson, & Sylvia Terbeck. 2016. Virtual morality: Transitioning from moral judgment to moral action? *PloS One* 11(10): e0164374.
- Fricker, Miranda. 2008. Scepticism and the genealogy of knowledge: Situating epistemology in time. *Philosophical Papers* 37(1): 27–50.
- Friedman, Jane. 2019. Inquiry and Belief. *Nous* 53: 296–315.
- Gerken, Mikkel. 2017. *On folk epistemology: How we think and talk about knowledge*. Oxford: Oxford University Press.
- Gilbert, Daniel T., & Timothy D. Wilson. 2007. Prospection: Experiencing the future. *Science*, 317: 1351–1354.

- Glöckner, Andreas, Benjamin Hilbig, & Marc Jekel. What is adaptive about adaptive decision making? A parallel constraint satisfaction account. *Cognition* 133(3): 641–666.
- Gold, Joshua I., & Michael N. Shadlen. 2007. The Neural Basis of Decision Making. *Annual Review of Neuroscience*, 30: 535–74.
- Greco, John. 2008. What's wrong with contextualism? *The Philosophical Quarterly* 58(232): 416–436.
- Green, Melanie C., & Timothy C. Brock. 2000. The role of transportation in the persuasiveness of public narratives. *Journal of personality and social psychology* 79(5): 701–721.
- Green, Melanie C., & John K. Donahue. 2009. Simulated Worlds: Transportation Into Narratives. In Keith D. Markman, William M. P. Klein, and Julie A. Suhr (eds), *Handbook of Imagination and Mental Simulation*, Routledge.
- Grice, H. P. 1989. *Studies in the Way of Words*. Harvard University Press.
- Halamish, Vered, & Nira Liberman. 2017. How much information to sample before making a decision? It's a matter of psychological distance. *Journal of Experimental Social Psychology* 71: 111–116.
- Hannon, Michael. 2013. The practical origins of epistemic contextualism. *Erkenntnis* 78(4): 899–919.
- Hansen, Nat. 2020. “Nobody would really talk that way!”: The critical project in contemporary ordinary language philosophy. *Synthese* 197: 2433–2464.
- Hansen, Nat, J. D. Porter, & Kathryn Francis. 2019. A corpus study of ‘know’: On the verification of philosophers' frequency claims about language. *Episteme* (First View).
- Hare, Brian. 2017. Survival of the friendliest: Homo sapiens evolved via selection for prosociality. *Annual Review of Psychology* 68: 155–186.
- Harman, Gilbert. 1986. *Change in View*. Cambridge, MA: MIT Press.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.
- Hawthorne, John, & Jason Stanley. 2008. Knowledge and Action. *Journal of Philosophy* 105(10): 571–590.
- Heitz, Richard. 2014. The speed-accuracy tradeoff: History, physiology, methodology, and behavior. *Frontiers in Neuroscience*, 8(150).
- Henderson, David. 2011. Gate-Keeping Contextualism. *Episteme*, 8(1): 83–98.
- Henderson, David, & John Greco (eds.). 2015. *Epistemic evaluation: Purposeful epistemology*. Oxford: Oxford University Press.
- Heyes, Cecilia M. 2018. *Cognitive Gadgets*. Cambridge, MA: Belknap Press.
- Holmes, Emily A., & Andrew Mathews. 2010. Mental imagery in emotion and emotional disorders. *Clinical psychology review* 30(3): 349–362.
- Holton, Richard. 2014. Intention as a model of belief. In Manuel Vargas & Gideon Yaffe (eds.), *Rational and social agency: Essays on the philosophy of Michael Bratman*. Oxford: Oxford University Press.
- Jackson, Alexander. 2012. Two Ways to Put Knowledge First. *The Australasian Journal of Philosophy* 90(2): 353–369.
- 2016. From Relative Truth to Finean Non-Factualism. *Synthese* 193(3): 971–989.
- 2019. Rampant Non-Factualism: A Metaphysical Framework and its Treatment of Vagueness. *Analytic Philosophy* 60(2): 79–108.

- Johnson, Dan R., Grace K. Cushman, Lauren A. Borden, & Madison S. McCune. 2013. Potentiating empathic growth: Generating imagery while reading fiction increases empathy and prosocial behavior. *Psychology of Aesthetics, Creativity, and the Arts* 7(3): 306–312.
- Kahneman, Daniel. 2003. A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58(9): 697–720.
- Kobayashi, Hiromi, & Shiro Kohshima. 1997. Unique morphology of the human eye. *Nature* 387/6635: 767–768.
- Kruglanski, Arie W., & Donna M. Webster. 1996. Motivated closing of the mind: “Seizing” and “Freezing”. *Psychological Review*, 103(2): 263–283.
- Ledgerwood, Alison. 2014. Evaluations in Their Social Context: Distance Regulates Consistency and Context Dependence. *Social and Personality Psychology Compass*, 8: 436–447.
- Lewis, David. 1996. Elusive Knowledge. *The Australasian Journal of Philosophy*, 74(4): 549–567.
- Liberman, Nira, & Yaacov Trope. 2014. Traversing psychological distance. *Trends in Cognitive Sciences* 18(7): 364–369.
- MacFarlane, John. 2014. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford University Press.
- McKenna, Robin. 2013. ‘Knowledge’ ascriptions, social roles and semantics. *Episteme* 10(4): 335–350.
- Mayseless, Ofra, & Arie W. Kruglanski. 1987. What makes you so sure? Effects of epistemic motivations on judgmental confidence. *Organizational Behavior and Human Decision Processes* 39(2): 162–183.
- Miller, Ryan, & Fiery Cushman. 2013. Aversive for Me, Wrong for You: First-person Behavioral Aversions Underlie the Moral Condemnation of Harm. *Social and Personality Psychology Compass* 7(10): 707–718.
- Mrkva, Kellen, Mark Travers, & Leaf Van Boven. 2018. Simulational fluency reduces feelings of psychological distance. *Journal of Experimental Psychology: General*, 147(3): 354–376.
- Nagel, Jennifer. 2010. Epistemic Anxiety and Adaptive Invariantism. *Philosophical Perspectives* 24: 407–435.
- . 2012. Mindreading in Gettier Cases and Sceptical Pressure Cases. In Brown & Gerken (eds), pp. 171–191.
- Newell, Ben, & Michael Lee. 2011. The right tool for the job? Comparing an evidence accumulation and a naive strategy selection model of decision making. *Journal of Behavioral Decision Making* 24(5): 456–481.
- Nook, Erik C., Jessica L. Schleider, & Leah H. Somerville. 2017. A linguistic signature of psychological distancing in emotion regulation. *Journal of Experimental Psychology: General* 146(3): 337–346.
- Open Science Collaboration. 2015. Estimating the Reproducibility of Psychological Science. *Science* 349: acc4716.
- Payne, John, & James Bettman. 2004. Walking with the scarecrow: The information-processing approach to decision research. In Derek Koehler & Nigel Harvey (eds), *Blackwell Handbook of Judgment and Decision Making*. Blackwell.

- Pearson, Joel, Thomas Naselaris, Emily A. Holmes, & Stephen M. Kosslyn. 2015. Mental imagery: functional mechanisms and clinical applications. *Trends in Cognitive Sciences* 19(10): 590–602.
- Peetz, Johanna, Anne E. Wilson, & Erin J. Strahan. 2009. So far away: The role of subjective temporal distance to future goals in motivation and behavior. *Social Cognition* 27(4): 475–495.
- Pinillos, N. Ángel. 2011. Some recent work in experimental epistemology. *Philosophy Compass* 6(10): 675–688.
- . 2012. Knowledge, experiments and practical interests. In Brown & Gerken (eds).
- Pinillos, N. Ángel, & Shawn Simpson. 2014. Experimental evidence supporting anti-intellectualism about knowledge. In James Beebe (ed.), *Advances in experimental epistemology*, London: Bloomsbury. Pp. 9–44.
- Powell, Derek, Zachary Horne, N. Ángel Pinillos, Keith J. Holyoak. 2015. A Bayesian framework for knowledge attribution: Evidence from semantic integration. *Cognition* 139: 92–104.
- Ratcliff, Roger, & Gail McKoon. 2008. The diffusion decision model: Theory and data for two-choice decision tasks. *Neural computation* 20(4): 873–922.
- Roets, Arne, Alain Van Hiel, Ilse Cornelis, & Barbara Soetens. 2008. Determinants of task performance and invested effort: A need for closure by relative cognitive capacity interaction analysis. *Personality and Social Psychology Bulletin* 34(6): 779–792.
- Rose, David, Edouard Machery, Stephen Stich, et al. 2019. Nothing at stake in knowledge. *Nous* 51(1): 224–247.
- Rysiew, Patrick. 2007. Speaking of Knowing. *Nous* 41(4): 627–662.
- Schilbach, Leonhard, Bert Timmermans, Vasudevi Reddy, Alan Costall, Gary Bente, Tobias Schlicht, & Kai Vogeley. 2013. Toward a second-person neuroscience. *Behavioral and Brain Sciences* 36: 393–462.
- Schwartz, Norbert. 2007. Attitude construction: Evaluation in context. *Social Cognition* 25: 638–656. (Special issue, *What is an attitude?*, ed. B. Gawronski.)
- Shea, Nicholas, Annika Boldt, Dan Bang, Nick Yeung, Cecilia Heyes, & Chris D. Frith. 2014. Supra-personal cognitive control and metacognition. *Trends in Cognitive Sciences* 18(4): 186–193.
- Simen, Patrick, David Contreras, Cara Buck, Peter Hu, Philip Holmes, & Jonathan D. Cohen. 2009. Reward rate optimization in two-alternative decision making: Empirical tests of theoretical predictions. *Journal of Experimental Psychology: Human Perception and Performance* 35(6): 1865–1897.
- Simon, Herbert. 1956. Rational choice and the structure of the environment. *Psychological Review* 63: 129–138.
- Smith, Eliot R., & Gün R. Semin. 2004. Socially situated cognition: Cognition in its social context. *Advances in experimental social psychology* 36, pp. 57–121.
- Soderberg, Courtney K., Shannon P. Callahan, Annie O. Kochersberger, Elinor Amit, & Alison Ledgerwood. 2015. The effects of psychological distance on abstraction: Two meta-analyses. *Psychological Bulletin* 141(3): 525–548.
- Söllner, Anke, Arndt Bröder, Andreas Glöckner, & Tilmann Betsch. 2014. Single-process versus multiple-strategy models of decision making: Evidence from an information intrusion paradigm. *Acta Psychologica* 146: 84–96.
- Sripada, Chandra Sekhar, & Jason Stanley. 2012. Empirical tests of interest-relative invariantism. *Episteme* 9(1): 3–26.

- Staffel, Julia. 2013. Can There be Reasoning with Degrees of Belief? *Synthese* 190(16): 3535–3551.
- 2019. How do beliefs simplify reasoning? *Nous* 53(4): 937–962.
- Standage, Dominic, Gunnar Blohm, & Michael Dorris. 2014. On the neural implementation of the speed-accuracy trade-off. *Frontiers in Neuroscience*, vol. 8, article 236.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford, UK: Oxford University Press.
- Sterelny, Kim. 2012. *The Evolved Apprentice*. MIT Press.
- Tetlock, Philip. 2002. Social Functionalist Frameworks for Judgment and Choice: Intuitive Politicians, Theologians, and Prosecutors. *Psychological Review* 109(3): 451–471.
- Tomasello, Michael. 2009. *Why We Cooperate*. MIT Press.
- Trautmann, Stephan T. 2019. Distance from a distance: the robustness of psychological distance effects. *Theory and Decision* 87: 1–15.
- Trope, Yaacov, & Nira Liberman. 2010. Construal-level theory of psychological distance. *Psychological Review* 117: 440–463.
- Turri, John. 2017. Epistemic contextualism: An idle hypothesis. *Australasian Journal of Philosophy* 95(1): 141–156.
- Turri, John, Wesley Buckwalter, & David Rose. 2016. Actionability judgments cause knowledge judgments. *Thought: A Journal of Philosophy* 5(3): 212–222.
- Van Boven, Leaf, & Eugene M. Caruso. 2015 The tripartite foundations of temporal psychological distance: Metaphors, ecology, and teleology. *Social and Personality Psychology Compass* 9(11): 593–605.
- Van Boven, Leaf, Joanne Kane, A. Peter McGraw, & Jeannette Dale. 2010. Feeling close: Emotional intensity reduces perceived psychological distance. *Journal of Personality and Social Psychology* 98(6): 872–885.
- Van Laer, Tom, Ko De Ruyter, Luca M. Visconti, & Martin Wetzels. 2014. The extended transportation-imagery model: A meta-analysis of the antecedents and consequences of consumers' narrative transportation. *Journal of Consumer Research* 40(5): 797–817.
- Wagenmakers, Eric-Jan. 2009. Methodological and empirical developments for the Ratcliff diffusion model of response times and accuracy. *European Journal of Cognitive Psychology* 21(5): 641–671.
- Warneken, Felix. 2015. Precocious prosociality: Why do young children help? *Child Development Perspectives* 9(1): 1–6.
- Weatherson, Brian. 2012. Knowledge, bets, and interests. In Brown & Gerken (eds).
- Weisberg, Jonathan. 2016. Belief in psyontology. *Philosophers' Imprint*.
- Wessler, Janet, & Jochim Hansen. 2016. The effect of psychological distance on automatic goal contagion. *Comprehensive Results in Social Psychology*, 1(1–3): 51–85.
- Wickelgren, Wayne A. 1977. Speed-accuracy tradeoff and information processing dynamics. *Acta Psychologica* 41(1): 67–85.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford, UK: Oxford University Press.
- 2005. Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge. *The Philosophical Quarterly* 55(219): 213–235.
- Yeh, Wenchi, & Lawrence Barsalou. 2006. The situated nature of concepts. *American Journal of Psychology* 119(3): 349–384.