

N. Ángel Pinillos

Arizona State University

1. Introduction

A common assumption in the epistemology literature is that ordinary knowledge attributions are sensitive to pragmatic factors (PFs). For example, it is assumed that people are less likely to attribute knowledge to an agent in some cases where what is at stake for the agent has increased significantly (even though the strength of her evidence or justification remained the same).¹ This empirical assumption is supported by philosophical reflection on ordinary usage (Fantl and McGrath 2009), experimental work (Pinillos 2012, Sripada and Stanley 2012, Pinillos and Simpson 2014, Francis et al. 2019) and reasonable theoretical assumptions about human cognition (Nagel 2008, Jackson, forthcoming).

This common assumption, however, has been challenged by a number of experimental philosophers. According to this surprising body of research, a certain type of straight-forward

* Penultimate draft

¹ Pragmatic encroachers make sense of this data is by positing that knowledge itself is sensitive to practical interests (Fantl and McGrath 2009). Competing accounts make sense of the data in different ways. For example, some contextualists claim that 'knows' is context sensitive and pragmatic factors play a role in setting the conversational standards which fix the content of 'knows' on a particular occasion DeRose (1995). Other theorists make sense of the data by appealing to psychological forces which distort our ordinary practices of attributing knowledge (Nagel 2008). What many defenders of these views have in common, however, is the acceptance of the common assumption that ordinary knowledge attributions are sensitive to pragmatic factors.

experimental paradigm *does not* reveal that knowledge attributions are sensitive to PFs (Buckwalter 2010, Buckwalter and Schaffer 2015, Feltz and Zarpentine 2010 May et al. 2010). In one of the most dramatic examples of this research program, a team of dozens of international researchers surveyed thousands of participants from around the world. Using the experimental paradigm I just mentioned, they discovered that throughout the world, ordinary knowledge attributions are, in general, not sensitive to what is at stake for the potential knower (Rose et al. 2019).

The upshot of this research is significant for epistemologists who care to make contact with our ordinary uses of epistemic notions. Specifically, researchers have thought that the common assumption in epistemology mentioned above supports certain theories which try to accommodate these pragmatic factors, including contextualism and interest relative invariantism (which we identify with pragmatic encroachment). Hence, if these experimental philosophers are right, this undercuts an important source of motivation for those theories. As Rose et al. say ‘If stakes really do not play any role in ordinary knowledge ascriptions, one of the main motivations for epistemic contextualism and interest-relative invariantism would be undermined’. Similarly, Schaffer and Knobe (2012, pg. 675) remark about these experimental studies: ‘it is easy to come away with the feeling that the whole contextualism debate was founded on a myth’. The experimental results, therefore, seem to have important ramifications for debates in epistemology.²

I want to argue that we should move away from this simple narrative about bank cases. In light of the study I present in this paper, I propose a more nuanced view about the lessons we should take away from these experimental studies. I distinguish two types of pragmatic factors that could

² Of course, there’s also an interest in how humans actually deploy epistemic concepts since we wish to have a better understanding of how people make epistemic assessments.

affect knowledge attributions: stakes and normative facts. For our purposes, a stakes factor is the downside of being wrong about the proposition you potentially know. A normative fact is the appropriateness of assuming the proposition you potentially know in your reasoning. There is a tight metaphysical relationship between stakes and normative facts. Uncontroversially, changes in the former can necessitate changes in the latter. Although simple guided reflection can establish this connection, philosophically naive subjects need not be aware of this intimate relation, or so I will argue. For these naive subjects, awareness of stakes need not spontaneously lead them to make a judgment about normative facts.

I think that from the perspective of folk epistemology, it is awareness of normative facts (not stakes) which can affect knowledge judgements. This is mainly why in many of the original bank studies (with ordinary subjects) no stakes effect was found. In contrast, philosophers display a stakes effect in those cases. This is because they tend to be aware of the intimate connection between stakes and normative facts. That transition is seamless for them. But at any rate, even for philosophers, it is the change in normative facts which explains their knowledge attributions in bank cases. As I will explain later, it's not like any change in stakes across vignettes is supposed to lead to a change in knowledge attributions. It is only a change in stakes to the extent that it leads to a change in normative facts which is supposed to lead to the asymmetric knowledge attributions. This suggest that it is the normative facts that are really doing the work.

In this paper I report on an original study which suggests that knowledge attributions by philosophically naive subjects are psychologically sensitive to normative facts (the effect, as we will see, is actually quite large). These normative facts are a species of pragmatic factors (PFs). Hence, the study suggests that ordinary attributions are sensitive to PFs. It is highly controversial that knowledge

itself is sensitive to these normative facts. If we are looking for theories of knowledge which straightforwardly reflect our ordinary practices, we would need to move away from traditional epistemic theories towards theories like contextualism or pragmatic encroachment which could better accommodate them. I endorse the conservative conclusion that our reaction to bank cases (suitably modified) do, after all, support those non-traditional epistemic theories.

I should add that although ordinary knowledge attributions by naive subjects are not psychologically sensitive to stakes, they are, in a metaphysical sense, sensitive to it. What do I mean by this? I mean that if you take a pair of cases which vary only in what is at stake, including elements that are necessitated by this difference, then ordinary people will attribute knowledge in different ways (so long as background facts are kept fixed, including the agents possessing the right amount of evidence). How does this happen? When I say that we vary only what is at stake including “including elements that are necessitated by this change”, we must include certain normative facts that are necessitated by this change. When we make these normative facts explicit in the vignettes, people will react to these normative facts and attribute knowledge in different ways across cases. To be clear, the psychological explanation for the behavior is their awareness of the normative facts. But metaphysically, it is the stakes that explains their behavior (plus background facts which are held fixed across cases). This is because the stakes difference necessitates and explains the difference in normative facts which in turn explain, psychologically, attributions of knowledge.

2. Bank Cases

So-called “bank cases” are used to motivate the thesis that attributions of knowledge are sensitive to pragmatic factors, in particular, what is at stake for the protagonist in the vignette (DeRose 2005). In one version, we are told about Bob, who wishes to deposit a check at his bank on a Friday night. Because lines are long, he decides to wait to the next day to deposit the check. His wife worries that the bank will not be open on Saturday, but Bob confidently reassures her that the bank will be open since he was there on a Saturday two weeks ago. In the low stakes case, we are told that it is not a big deal if the check is not deposited by Saturday. In the high stakes case, we are told that it is big deal if the check is not deposited by Saturday. In the particular version we are investigating (which comes from the international studies reported in Rose et al. 2017), Bob asserts that he knows the bank will be open on Saturday (in both the low and high stakes cases alike). Professional philosophers presumably tend to judge that Bob’s assertion is true in the low stakes case, but false in the high stakes case—thus revealing that these attributions of knowledge are sensitive to PFs.

However, this asymmetry in responses to these and related bank cases has not been found in the general population. As mentioned above, some experimental philosophers cite this result as a challenge to the epistemologist’s assumption that ordinary attributions of knowledge are sensitive to pragmatic factors, and stakes in particular.

I end this section by displaying the traditional bank cases in full. This is the version that appeared in the international study from Rose et al. 2017.

Low Stakes

Bob and his wife are driving home on a Friday afternoon. They both received some money earlier in the day, and so they plan to stop at the bank on the way home to deposit it. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Although they generally like to deposit any money they receive at the bank as soon as possible, it is not especially important in this case that it be deposited right away, and so Bob suggests that they drive straight home and deposit their money on Saturday morning. His wife says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." Bob replies, "No, I know the bank will be open. I was just there two weeks ago on Saturday. It was open until noon." As a matter of fact, the bank will be open on Saturday morning.

High Stakes

Bob and his wife are driving home on a Friday afternoon. They both received some money earlier in the day and so they plan to stop at the bank on the way home to deposit it. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. They have recently written a very large and very important check. If the money is not deposited into their bank account before Monday morning, the important check they wrote will not be accepted by the bank, leaving them in a very bad situation. Bob suggests that they drive straight home and

deposit their money on Saturday morning. His wife says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays.” Bob replies, “No, I know it’ll be open. I was just there two weeks ago on Saturday. It was open until noon.” As a matter of fact, the bank will be open on Saturday morning.

3. Modifying the Bank Cases

I want to argue that there’s a big difference between how philosophers and ordinary people interpret traditional bank cases. Once we modify the traditional bank cases to bridge this gap, ordinary people will respond to the bank cases the way philosophers do. To see this, first note that the difference between the High and Low stakes bank cases are the costs associated with being wrong about the bank being open on Saturdays. This is certainly a difference in pragmatic factors or what is at stake. But no philosopher holds the view that *any* change in what is at stake automatically leads to a change in ordinary knowledge attributions.

What is it about bank cases which leads philosophers to think that there would be a difference in knowledge facts across the cases? I think that philosophers naturally interpret the difference between High and Low Stakes as involving a difference in normative facts—only in High Stakes is Bob engaged in subpar reasoning. Specifically, in Low Stakes, it’s ok for Bob to assume that the bank will be open on Saturday given his evidence (his evidence is the fact that he visited the bank a couple of Saturdays ago and it was open). But in High Stakes, it’s not ok for Bob to assume that the bank will be open on Saturday given his evidence. Let’s call this interpretation of the bank cases where part of fully understanding the vignettes involves an assessment about what is ok for Bob to

assume, 'the normative interpretation'. Let's call true (in the story) claims like 'It's ok for Bob to assume that the bank will be open on Saturday', the 'normative facts'.

Philosophers are first introduced to bank cases in the context associated with reading a philosophy journal article or taking a philosophy class. The bank cases are understood to be part of a larger investigation into epistemic matters. They often arise to motivate the connection between knowledge attributions and pragmatic factors, including the normative status of actions. In particular, they come up in the pragmatic encroachment literature where facts about the rationality of action are bound up with knowledge. Philosophers take our intuitive reaction to cases similar to the bank vignettes to support principles connecting the rationality of action and knowledge. Given all this, it shouldn't be surprising to learn that philosophers adopt the normative interpretation in bank cases. Remarks by Fantl and McGrath (2009b) suggest this much. They say:

...why, then, do subjects say 'I don't know' in cases like Bank Case B [a high stakes case]? One explanation is that in such cases it is clear that it takes knowledge that p in order to properly act on p ' (footnote 14).

Fantl and McGrath say that in the high stakes cases, it is *clear* that it takes knowledge that p in order to properly act on p . The assumption here is that subjects must interpret these high stakes as being ones where it is *not* proper to act on the premise that the bank will be open on Saturday. That's why they judge that the agent fails to know. This assumption is not explicitly stated in the vignettes. It is just the normative interpretation.

Fantl and McGrath also discuss another famous high stakes case (Cohen 1999). In this scenario, we are told about some protagonists in a high stakes case who think they do not know that a flight stops in Chicago (based on an itinerary printout they got from a travel agent). Fantl and McGrath say that ‘the case is devised so that they [the protagonists] should check with the information desk, given all that is at stake.’ That is, the case is designed to be interpreted in such a way that the protagonists should not assume that the flight stops over in Chicago without first gathering more evidence. In short, the case is ‘devised’ by Cohen to trigger the normative interpretation (even though the normative claim is not explicitly specified in the vignettes).

There’s good reason to think, therefore, that when philosophers read the bank vignettes in philosophical contexts, they are more likely to adopt the normative interpretation. In contrast, participants who take surveys are not philosophers nor are they aware of the surrounding debate and literature associated with the bank cases. They are less likely to read between the lines and adopt the normative interpretation. If this is right, when bank cases are presented to non-philosophers, we must ensure participants lock on the normative interpretation. I devised modified bank cases which help agents do this. These are just like the bank cases I above except that I make a small addition. The change is italicized below (in the actual materials given to participants nothing was italicized).

Modified Low Stakes

Bob and his wife are driving home on a Friday afternoon. They both received some money earlier in the day, and so they plan to stop at the bank on the way home to deposit it. But as they drive past the

bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Although they generally like to deposit any money they receive at the bank as soon as possible, it is not especially important in this case that it be deposited right away. Bob suggests that they drive straight home and deposit their money on Saturday morning. His wife says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." Bob replies, "No, I know the bank will be open. I was just there two weeks ago on Saturday. It was open until noon." As a matter of fact, the bank will be open on Saturday morning. *The stakes are low for Bob and his wife. So based on Bob's evidence, it's appropriate for them to assume in their reasoning that the bank will be open on Saturday.*

Modified High Stakes

Bob and his wife are driving home on a Friday afternoon. They both received some money earlier in the day, and so they plan to stop at the bank on the way home to deposit it. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. They have recently written a very large and very important check. If the money is not deposited into their bank account before Monday morning, the important check they wrote will not be accepted by the bank, leaving them in a very bad situation. Bob suggests that they drive straight home and deposit their money on Saturday morning. His wife says, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." Bob replies, "No, I know the bank will be open. I was just there two weeks ago on Saturday. It was open until noon." As a matter of fact, the bank will be open on

Saturday morning. *The stakes are high for Bob and his wife. So based on Bob's evidence, it's not appropriate for them to assume in their reasoning that the bank will be open on Saturday.*

In this section I argued that since no philosopher thinks any change in stakes should lead a change in knowledge attributions, something else must be thought to be responsible for the change in knowledge attributions in the classic bank cases. This is likely to be the normative facts and so I think philosophers adopt the normative interpretation of the classic bank cases. Accordingly, I developed new bank cases which make the normative interpretation explicit. My prediction is that when this interpretation is made explicit, participants will respond to these modified bank cases the way philosophers react to the traditional bank cases.

4. Modified Bank Cases Differ in Stakes

I argued above that Modified High and Low Stakes cases are explicit renditions reflecting how bank cases are interpreted by philosophers. But I could be wrong about this sociological fact. It's not essential to the main point of this paper. What is of greater philosophical significance is whether (a) ordinary people really ascribe knowledge differently across the Modified High and Low Stakes cases, and whether (b) the Modified High and Low cases only differ only in pragmatic factors or stakes (which one exactly will be important). I will show (a) in the next section. But I now discuss (b).

When what is at stake changes across scenarios and everything else remains the same as much as possible, including the strength of evidence held by the agent in question, then we say the scenarios differ *only* in stakes. In the original (unmodified) bank cases, the scenarios differ only in

stakes. Now, a pair of vignettes may differ more broadly with respect to pragmatic factors. If two stories differ only in stakes, then they differ with respect to pragmatic factors. But not necessarily the other way around. For example, my former student Joseph Shin (2014) showed that the length of time one has to make a decision affects attributions of knowledge (keeping the evidence the same as well as the actual time the agent takes to make the decision). I take this to be a difference in pragmatic factors that don't involve a difference in stakes.

In the modified bank cases, Bob's evidence for the proposition that the bank will be open on Saturday presumably remains the same across both conditions. What is at stake changes as do the normative facts. In Modified Low Stakes, it is stipulated that it is appropriate for Bob to assume that the bank will be open on Saturday. In Modified High Stakes, it is stipulated that it is not appropriate. Are these difference in normative facts amount to a difference in pragmatic factors? Are they differences that follow from the difference in stakes?

Here is the thesis that I will defend about the differences between the Modified High and Low Stakes Cases:

MINIMAL: The Modified High and Low Stakes cases differ only in stakes

I will argue MINIMAL is true in 4.1. And in section 5, I will show that ordinary knowledge attributions differ greatly between the Modified Low and High Stakes conditions. This means that there's a sense in which ordinary knowledge attributions are sensitive to stakes. However, it doesn't follow that what caused the differing knowledge attributions across the conditions for ordinary agents is an awareness of stakes. It may be that what is responsible for the asymmetric knowledge

attributions are elements which are downstream consequences of the stakes—which are now explicitly stated in the stories. In our case, these downstream explicitly laid out elements are the normative facts—whether it’s ok for our agents to assume that the bank will be open on Saturday.

We will describe the sense in which knowledge attributions are sensitive to stakes as ‘metaphysical’. The sense we just discussed in which knowledge attributions are sensitive to normative facts is ‘psychological’. I will argue for the following two theses:

METAPHYSICAL STAKES: Ordinary knowledge attributions are sensitive to stakes in the sense that ordinary knowledge attributions vary across a pair of cases which differ only in stakes, but where the psychological explanation for this variation in knowledge attributions may be due to something else which is metaphysically necessitated by the difference in stakes (this “something else” is explicitly mentioned in the cases).

PSYCHOLOGICAL NORM: Ordinary knowledge attributions are sensitive to normative facts in a psychological-causal sense.

An important question is whether the normative facts explicitly mentioned in the modified cases count as pragmatic factors (as opposed to traditional epistemic factors). I will argue in section 9 that the answer is ‘yes’:

PRAGMATIC NORM: The explicitly mentioned normative facts in the modified bank cases (whether or not it's ok for Bob to assume that the bank will be open on Saturday) is a pragmatic factor, as opposed to a traditional epistemic factor.

As I mentioned above, there's also the question of whether stakes play a psychological role in ordinary knowledge ascriptions. I discuss this later on, but it is likely that it does not, at least on a standard interpretation of 'psychological'. For philosophers who think about the bank cases, however, it is likely that stakes play a psychological role in their knowledge judgments. We will discuss these claims further. But we first need to argue for MINIMAL. MINIMAL plus the results of an experiment I report in section 5 yield METAPHYSICAL STAKES.

4.1 Argument for MINIMAL

MINIMAL: The Modified High and Low Stakes cases differ only in stakes

To see how the argument for MINIMAL works, first note that when we truthfully say that two cases differ *only* in some feature F , we allow that the cases may differ with respect to some other feature G which is necessitated by this difference in F . Here are some examples. Imagine a man walking down

the street carrying three tomatoes and a head of lettuce. Now modify this story so that the man is now holding four tomatoes and a head of lettuce. I think we can correctly say that the stories differ *only* in that the man is carrying three tomatoes in one but four tomatoes in the other story. But I think we can also truthfully say, without contradicting ourselves, that one difference between the stories is that the man is holding an odd number of tomatoes in one scenario but not the other. We allow this way of talking because the first difference necessitates the second difference.

Here's another example. We can correctly say of two stories featuring a long rod that they differ *only* in the length of the rod (not the other dimensions) while also correctly saying without contradicting ourselves that they also differ with respect to the volume of said rod. This is because, necessarily, a difference in length requires a difference in volume when all other dimensions remain the same.

And yet another example. I can correctly say that two scenarios that differ only with respect to the mass of an object that they also differ with respect to the weight of the object (again, without contradiction). Note that I am not assuming that a difference in mass always requires a difference in weight (in a setting with zero gravity the entailment does not hold). But supposing that it is implicitly assumed that the scenario takes place on earth and that current physical laws are in place, then we could correctly say of scenarios which only differ in the mass of an object that they also differ with respect to the weight of said object. This is because a difference in mass necessitates a difference in weight given certain background assumptions which are taken on board in the scenarios.

The point extends to normative facts. Suppose I correctly describe a pair of scenarios as differing only in that in the first scenario, an agent is in great physical pain. I think we can also say that the scenarios differ in some normative fact. For example, that the agent is worse off in the first

scenario than in the second scenario. This is compatible with the original claim that the scenarios differ *only* with the amount of pain the protagonist feels. The difference in normative facts are necessitated by the difference in pain facts.

Finally, let us consider the Modified Bank Cases. I think we can describe them as differing *only* in stakes. But this difference, as it appears in the bank cases and given some background assumptions taken on by philosophers, necessitates a difference in what is acceptable for Bob and his wife to assume in their reasoning. Hence the cases also differ in normative facts (and so MINIMAL is true).

Let's examine this idea. Something that is acceptable to assume in a setting will not be acceptable to assume in a different setting where the stakes are raised sufficiently high (leaving everything else the same as much as possible).³ For example, suppose it's acceptable to assume in my reasoning concerning party planning that you will show up at my house in the next two hours. I base my assumption on what you said three weeks ago. But now imagine that my reasoning concerns an important ransom I need to pay to save my life. If you don't show up to my house soon, I will die. The stakes are really high. In this circumstance, it's no longer acceptable for me to assume that you will show up in two hours based on what you said three weeks ago. I need to call you to figure out when exactly you will be coming. This example illustrates how a difference in stakes necessitates a difference in normative facts.

Think about it another way. Assuming a premise in reasoning is a type of *action*. The acceptability of an action depends on the downside of taking the action.⁴ Consider a situation in

³ Unless, perhaps, the justification is maximal.

⁴ I am assuming throughout that the agent is aware of the downside or the stakes.

which it is acceptable to take an action. There's a level at which you can raise the downside of taking the action, leaving everything else the same, such that the action becomes no longer rational—if you raise the stakes enough, the acceptable action becomes not acceptable. In short, there are cases where a stakes difference necessitates a difference in normative facts.

Now, I am not saying that *any* change in facts about what is at stake will lead to a change of normative facts. Just that there are cases in which a change in only the stakes leads to a change in what is acceptable to assume. With respect to the traditional bank cases, what I am saying is that there's a natural interpretation of the bank cases in philosophical settings where the difference in stakes between them is large enough that it necessitates a difference in normative facts (what is acceptable to assume).

With respect to modified bank cases, I think the idea that the stakes difference in the modified bank cases necessitates a difference in normative facts is also true. Consider how the normative facts are described in the modified cases. They are described as being the result of stakes: *The stakes are high for Bob and his wife. So based on Bob's evidence, it's not appropriate for them to assume in their reasoning that the bank will be open on Saturday.* This suggests that MINIMAL is true.

The idea that sometimes a stakes difference leads to a difference in normative facts is not a controversial claim. What *is* controversial is that there's an interesting connection between normative facts and knowledge. What this means for us is that MINIMAL should not be something that is denied by parties in the debate surrounding knowledge and stakes (in virtue of them holding their positions). In short, accepting MINIMAL is not tantamount to begging the question against people in the stakes and knowledge debate.

Finally, I think it should be noted that this metaphysical entailment between stakes and normative facts is not something ordinary agents spontaneously notice. I think it requires reflection, perhaps going carefully through something like the reasoning I went through above. Thus, we should not expect simple surveys to be able to elicit this inference. Of course, we can imagine further experimental work where you guide participants to appreciate the connection between stakes and what is rational to assume. It would be interesting to see how such agents react to the traditional Bank cases.

5. The Study

Materials

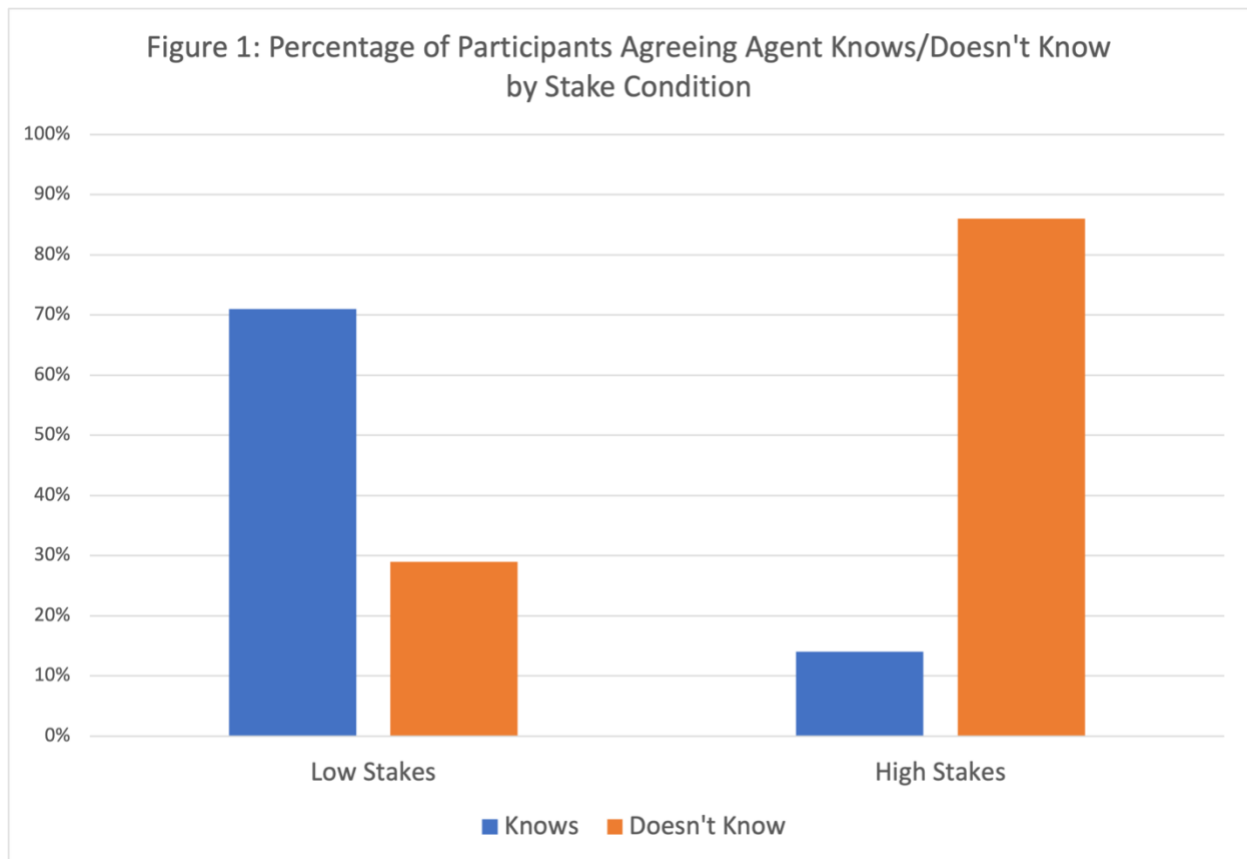
I conducted a study to test whether ordinary attributions of knowledge vary across Modified High Stakes and Modified Low Stakes. 100 subjects were recruited through Amazon Turk to fill out an 8 minute survey. Participants were all Amazon Master worker adults residing in the US. They were compensated \$1.50 each.

Each participant was assigned to either the Modified High Stakes or Modified Low Stakes (see above). After reading the study, they were asked two reading comprehension questions: ‘It is very important that Bob and his wife deposit the check right away’ and ‘it is appropriate for Bob and his wife to assume that the bank will be open Saturday’. Each of these questions were followed by three possible answers ‘True’, ‘False’ and ‘Not Applicable’. The third question gets at knowledge attributions. I deploy the same type of probe used by Rose et al. 2017:

In your personal opinion, which better describes Bob's situation?

(a) Bob knows that the bank will be open on Saturday

(b) Bob thinks he knows that the bank will be open on Saturday, but he doesn't actually know the bank will be open on Saturday.



Results

I disqualified 11 participants for failing at least one of the comprehension checks. As we will see, the disqualification of these individuals does not affect the main results of the study. Here are the main findings. in Modified Low Stakes, 35 (71%) people agreed Bob knows while 14 (29%) people

said he didn't know. In Modified High Stakes, only 5 (14%) people agreed that Bob knows the bank will be open on Saturday while 35 (86%) agree that he doesn't know. These statistics are depicted in figure 1.

I performed a chi-square test to determine the independence of knowledge attributions and stakes. The relationship between the variables was significant $X^2(1, N = 89) = 30.9059, p < .01$. Cramer's V is .58, revealing a large effect size.⁵ As a point of comparison, perhaps the most celebrated result in experimental philosophy, the "Knobe Effect" displayed a similar large effect size when the original data was collected (Knobe 2003): $V = .59$.

What do these results tell us? Given that the only difference between the cases are the stakes (Minimal), then the results tell us that ordinary knowledge attributions are sensitive to stakes (METAPHYSICAL STAKES). But suppose that we are interested in the psychological cause of the difference between the knowledge attributions? Then it is plausible that this difference is caused by something downstream from stakes. It is likely that it is caused by the different normative descriptions of the cases which are now made explicit. I say this is plausible because (as we've been emphasizing) in the unmodified bank cases, where the normative descriptions are not made explicit, no stakes effect was found. And now when the normative descriptions are made explicit, we found a large effect. Hence, it is likely that our study supports PSYCHOLOGICAL NORM.

⁵ If we include all the disqualified participants, we get a similar result $X^2(1, N = 100) = 21.1685, p < .001, V = .46$

This result also supports the idea that ordinary knowledge attributions are sensitive to pragmatic factors in a psychological sense. As I will argue later, these normative facts are just pragmatic facts—they are facts about what course of action you should take (PRAGMATIC NORM). These normative facts are not ‘truth conducive’ elements traditionally tied to knowledge. So if we find that ordinary knowledge attributions are sensitive to these normative facts, we vindicate the spirit if not the letter of the common assumption that was thought to be challenged by experimental philosophers. More significantly, insofar as pragmatic encroachment and contextualist theories get support from the idea that ordinary attributions of knowledge are sensitive to pragmatic factors like the normative facts, then any support for those theories which was thought lost is regained.

6. What went wrong with the original bank cases?

So what went wrong with the original bank cases? Why don't non-philosophers ascribe knowledge the way philosophers do? I already suggested a possible answer. Since ordinary people are not explicitly told whether Bob's premise that the bank will be open on Saturday is appropriate (the normative interpretation philosophers lock onto), then they do not attend to those normative facts. If they were to attend to those facts, then their judgements would align with the judgments philosophers make in bank cases. This much is suggested by the study I reported above.

There are other explanations in the literature about why stakes don't seem to affect ordinary knowledge ascriptions in the bank cases. First, the simple view is that stakes don't play a role in knowledge ascriptions in general. Our discussion above suggests this option needs to be clarified. If it is claimed that ordinary knowledge ascriptions are not sensitive to stakes in the metaphysical sense, then I think our study and accompanying discussion refutes this thesis. If it is meant that ordinary

attributions are not sensitive to stakes in a psychological sense, then it may be correct (though the thesis requires some further clarification about what we mean by ‘psychological’—I say more about this at the end of the paper). Finally, if the explanation is that ordinary knowledge attributions are not sensitive to pragmatic factors more generally (including those that are downstream from a difference in stakes) then this thesis is also refuted by the study (see the theses of PSYCHOLOGICAL NORM and PRAGMATIC NORM).

A second explanation for why no stakes effect was detected in the unmodified bank cases is that, as Keith DeRose (2011) argued, there’s a key feature of High Stakes which could create a confound. To see this, we have to assume DeRose’s contextualist treatment of ‘knows’. He has famously held that knowledge attributions are context sensitive expressions which express different propositions reflecting different epistemic standards depending on the conversation context at play. For example, in a low standards context, the sentence ‘I know p ’ may express a true proposition. But in a high standards context, that sentence may express a different proposition that is false. There is no fully worked out theory about which factors of a conversation affects the epistemic standards. But higher stakes are thought to trigger higher standards. This means that the null experimental findings on the unmodified bank cases put some pressure on contextualist theorists.

In response to this experimental challenge, DeRose points out that there are other contextual features besides stakes which can move epistemic standards. For example, if an agent asserts ‘I know p ’, this speech is “accommodated”. In certain cases, this may move the standards down so as to make that speech true. This means that an assertion of ‘I know p ’ in a high stakes context may counter-act the standard-raising effects of a high stakes setting which would otherwise predict that the speech is

false. This then gives us an explanation why the stakes effect was not found in the traditional bank cases.

A third account of why the original bank cases were not able to elicit a stakes response is due to Pinillos (2012). Just like DeRose, I also focused on the fact that in the High Stakes cases, the protagonists claim to ‘know the bank will be open Saturday’. Let’s grant that higher stakes in this kind of case means that Bob is less likely to know (we are imagining a response on behalf of the pragmatic encroacher). Still, the fact that Bob *says* he knows is good reason to think he actually knows. After all, he is in a much better position than us, as readers, to judge whether he knows the bank will be open on Saturday. Let’s add more detail to this idea.

We are told in both the unmodified High and Low Stakes cases that the evidence Bob has for the proposition that the bank is open on Saturday is that he was there two weeks ago on a Saturday. However, participants may not think of the quality of Bob’s evidence in the same way across both High and Low Stake cases. Suppose, for example, that it is part of folk epistemology that stakes affects knowledge attributions (again, we are imagining a response on behalf of the pragmatic encroacher). Participants may react to the unmodified High Stakes cases by trying to make sense of how Bob may both (a) know that the bank is open on Saturday (because he says he knows this) and also (b), be in a high stakes situation. They can make sense of these things by construing the evidence possessed by Bob as being of higher quality. Although we are told that Bob’s evidence for the proposition that the bank will be open on Saturday was that he was there two weeks ago on a Saturday, the overall description of the evidence can be interpreted in different ways leading to different assessments on the quality of the evidence. For example, we can fill in different details about how likely it is that the bank would have changed their hours in the last two weeks. This background evidence can affect

the felt probative value of the key piece of evidence (that Bob was at the bank two weeks ago on a Saturday). What determines our estimation of how likely it is that the bank will have changed its hours (or how likely Bob thinks the bank will have changed their hours)? This estimation can increase as stakes increase as we try to make sense of how Bob can confidently assert he knows in a high stakes situation.

If I am right that in the unmodified bank cases, the quality of the evidence is not construed to be the same across the conditions, then we have a possible confound. This is because the Bank cases are supposed to differ only in what is at stake, not in the evidence possessed by the protagonist. This would explain why participants don't note a difference in knowledge facts across the conditions (despite the difference in stakes).

Using this idea as background, I developed a type of experimental paradigm I dubbed 'Evidence-Seeking'. This kind of probe uses both high and low stakes vignette, but instead of asking participants whether an agent's claim to 'know' is true, we ask them how much evidence the protagonist would need to gather before he can come to know some claim is true. In one experiment, I described a student who is proofreading his paper to try to correct the typos. In one condition, the stakes are high—if there's a single typo, his grade will fall below a certain important threshold. In the conditions where the stakes are low, it doesn't matter much if there's a typo in the paper. When I asked participants to tell me the number of times the student should proofread the paper before he knows there are no typos, I found a significant difference between responses to the low and high stakes versions of the probe. On average, participants in the low stakes condition said that the student has to proofread his paper twice before he knows the paper has no typos. But in high stakes, that number was 5. I took these results to support the thesis that ordinary knowledge attributions are

sensitive to stakes. I argued that the problem with the traditional bank cases is that they don't keep the evidence constant across the scenarios. The 'evidence-seeking' probes get around this problem. This experimental paradigm has been replicated (Buckwalter and Schaffer 2015, Francis et al. 2019).

Alex Jackson (forthcoming) provides yet a fourth account of what may have gone wrong in the unmodified bank cases. He argues that for stakes to affect our knowledge ascriptions, participants should be "psychologically close" to the protagonists in the vignette. But in fact, the stories are too remote and abstract for the participants themselves to "calibrate" to the stakes at play for the protagonist. Now, to explain why philosophers' reaction to bank cases are sensitive to stakes, he says that philosophers are psychologically close to the protagonists. The reason this happens is that philosophers are more likely to form rich imagery when reading the vignettes—which induces psychological closeness.

We just surveyed four possible explanations for why no stakes effect on ordinary people was found in traditional bank cases. These accounts would need to be modified to account for the new data. Recall that according to this new data, when we make explicit whether or not it is appropriate for Bob to assume that the bank will be open on Saturday (the normative claim), participants decidedly make attributions that are sensitive to stakes/normative facts.

DeRose's contextualist explanation would presumably need to say how making the normative claim explicit raises or lowers the standards for knowledge. And my own prior account faces the challenge of saying why evidence is now kept the same in both the modified Low and High Bank Stakes cases (where presumably it wasn't in the original bank cases). And Jackson needs to explain why the effect is detected in the modified cases. It seems unlikely that adding the one normative

claim would all of a sudden make the protagonists in the vignettes psychologically close to the participants.

Further work is needed to get to the bottom of why the original bank cases did not elicit the stakes effect on ordinary people. And this goes for the account I propose here—that it is attention to the normative facts which psychologically explains the effect we detected. In the next section I provides some of this detail.

7. Explaining Our Reaction to Bank Cases.

I've been arguing that attending to the normative facts in the modified bank cases (whether or not it is appropriate for Bob to assume that the bank will be open Saturday) causes participants to make asymmetric knowledge attributions. This is PSYCHOLOGICAL NORM (Ordinary knowledge attributions are sensitive to Normative facts in a psychological sense). But what exactly is the link between a person attending to the normative facts concerning the bank stories and their attributions of knowledge? Many philosophers have posited an intimate connection between these notions. One example is the Knowledge-Action principle (Hawthorne and Stanley 2008): Treat the proposition p as a reason for acting only if you know p . Philosophers have even argued for strengthened bi-conditional versions of this principle.

Such principles, however, cannot explain why *attending* to normative facts predicts making any type of knowledge attributions. These principles say nothing about what happens when an agent attends to certain normative facts or when they should attribute knowledge. They usually just talk about the connection between normative facts and knowing itself. However, such principles can be

transformed to make those predictions if we think of them as part of our folk epistemology as opposed to some deep metaphysical truth hidden in Plato's heaven (Pinillos 2019). That is, we can think of them as principles which are explicitly or implicitly accepted by ordinary agents which play a role in explaining our ordinary epistemic attributions and practices. For example, if ordinary people implicitly or explicitly accept something like the Knowledge-Action principle, then this would explain why them attending to the claim that it is acceptable for Bob to use the proposition that the bank will be open on Saturday as a premise in reasoning (in a low stakes case) leads them to judge that Bob knows that the bank will be open on Saturday. A strengthened principle (to a bi-conditional) can be similarly invoked to explain their reaction to the high stakes cases.

There is much to recommend the idea that principles similar to Knowledge-Action are accepted by ordinary agents (and are part of folk epistemology). It can explain why we ordinarily use 'knows' to blame and praise. For example, we say things like 'you should have worn the tie, you knew the president would be here' or 'you didn't know the president would take the day off, you should have worn a tie.' See Fantl and McGrath (2009) for more details.

Let's streamline our explanatory sketch of the human reaction to the modified bank cases. First, ordinary agents attend to the normative description in the modified bank cases. That is, in the low stakes probe they attend to the fact that it is appropriate for Bob to assume that the bank will be open on Saturday. In the high stakes probe they attend to the fact that it is not appropriate for Bob to assume that the bank will be open on Saturday. Next, they appeal to some implicitly or explicitly held knowledge-action link principle to then make a judgment on whether Bob knows that the bank will be open on Saturday.

In the unmodified bank cases, since the normative facts are not made explicit in the story and ordinary agents do not adopt the normative interpretation, they do not make different knowledge attributions across the cases. But what about philosophers? Their knowledge attributions shift between stakes in the unmodified bank cases. Why is that? I suggested earlier that it is because they adopt the ‘normative interpretation’. That is, they interpret the Low Stakes case as being one in which it is appropriate for Bob to assume that the bank will be open on Saturday, and they interpret the High Stakes case as being one in which it is not appropriate for Bob to assume that the bank will be open on Saturday.

Why should we think that philosophers as opposed to ordinary people (taking surveys) are more likely to adopt the normative interpretation? There are a few reasons. First, it explains the results of the study (modified bank cases): once we force the normative interpretation on ordinary agents, they respond to bank cases the way philosophers do. Second, the context in which philosophers read the bank vignettes is very different from the context in which survey participants read those vignettes. The former is more likely to trigger the normative interpretation. I explained why earlier in the paper.

Two further points are relevant. First, when philosophers read about the bank cases, they are exposed to both the high and low stakes versions at once. This induces them to compare and contrast the cases—to look for differences that stand out. One such difference is the difference in stakes and everything else that follows from this difference (given the background assumptions) including the normative facts (see the argument for Minimal). In contrast, each participant who takes a survey is usually only given one bank cases. These participants would seem to be less likely to identify the key

changes between the stories including those, like the normative facts, that are inferred from the explicitly laid out differences.

Second, recall that in the original bank cases we discussed, both high and low stakes conditions feature the protagonist asserting that he knows that the bank will be open on Saturday. In contrast, in DeRose's original bank case (which philosophers are familiar with), the protagonist in the high stakes condition asserts that he *doesn't know* that the bank will be open. As I mentioned above, this is an important difference between the cases which could further contribute to the difference in reaction between philosophers and non-philosophers.

8. So What Role do Stakes Play?

I argued that in the modified bank cases, what explains, psychologically, the difference in knowledge attributions across High and Low stakes is the difference in normative facts. In Low Stakes, agents infer from the fact that it is appropriate for Bob to assume that the bank will be open on Saturday to the claim that he knows this. In High stakes, agents infer from the fact that it is not appropriate for Bob to assume that the bank will be open on Saturday to the claim that he doesn't know this. But what role do the stakes play in the explanation of these knowledge judgments?

It may very well be that in the modified bank cases, stakes play no cognitive role for study participants in the sense that agents don't infer from stakes facts to normative facts. They start from normative facts and infer the knowledge facts. This possibility is consistent with ordinary knowledge ascriptions being sensitive to stakes in the metaphysical sense (see section 4). It is also consistent with

the possibility that stakes play a psychological role in knowledge ascriptions in some other less standard sense of ‘psychological’ as I explain below.

As I argued earlier, the normative facts are necessitated by the stakes facts (see the argument for Minimal). I don’t take this to be a controversial thesis. It is an inference whose validity can be discerned through philosophical reflection. But it doesn’t follow that this is an inference that is readily made by ordinary subjects. And it doesn’t follow that the inference is part of our folk epistemology. If it were, then we would expect ordinary knowledge attributions to be sensitive to stakes in a stronger psychological sense. But they are not, as the null results in the original bank experiments suggest. Instead, if we are interested in psychological explanations, we should say that it is attention to the normative facts which predict knowledge attributions, at least for bank cases.

Here’s a schematic causal representation of the relationship between stakes, normative facts, and knowledge attributions:

Stakes \longrightarrow Normative \implies Knowledge

The single arrow ‘ \longrightarrow ’ represents an inference that is made through philosophical reflection. The double arrow ‘ \implies ’ represents an inference of folk epistemology. This is the picture that emerges from the fact that the stakes effect is not detected in the traditional bank cases when those cases are presented to ordinary people, but it is detected in the modified bank cases.

There’s some existing research that backs up this general point. John Turri, Wesley Buckwalter and David Rose (2016) argue that actionability judgments cause knowledge judgments. For example, in one study, participants were presented with either a high stakes or a low stakes case.

Both cases involve an intelligence officer, Jennifer, filing a report on Ivan, a foreign operative. In the low stakes case, Jennifer hears from a source that Ivan has stopped his low carb diet and has quit jogging. There's little at stake if she messes up this report. In the high stakes case, Jennifer hears from a source that Ivan is selling weapons to a terrorist and is a threat. There's a lot at stake if she messes up this report. Participants were asked a number of questions including an "actionability" question—participants were asked to rate their agreement with this statement: Jennifer should write in the report that Ivan no longer diets/is a threat. They were also asked a knowledge question. They were asked to rate their agreement with this statement: Jennifer knows that Ivan (no longer jogs/is a threat).

Actionability is a normative notion. When participants are asked the extent they agree with the claim that Jennifer should write in the report that such and such, they are asked to rate a normative claim. Now to be sure, this type normative claim is not the same as the normative claim in the modified bank cases. The latter were designed to correspond to those in the various knowledge-action principles in which acceptability of a premise is at issue. For Turri et al. the normative claim concerns the appropriateness of writing a report. This adds a bit of complication since it may obviously be acceptable to write something down which is not acceptable to assume. And also, it may be acceptable to assume something without it being acceptable to write it down. But we will put this issue aside.

Here are some of their results. First, in one study, knowledge attributions were detected to be sensitive to stakes. Second, using the Greedy Equivalence Search algorithm, they argued that actionability judgments directly caused the knowledge judgments. In a second study, they similarly found that actionability judgments directly caused knowledge judgments. However, they did not

detect that knowledge attributions were sensitive to stakes. These results support our general model that normative facts psychologically explain knowledge attributions. It is also consistent with the idea that the psychological role of stakes is much less clear.

It is worth pointing out that we should resist describing the situation with the modified bank cases as one in which the Normative Facts mediate the effect of stakes on knowledge. This would be a misleading way of putting it. There is no reason to think that (at least in the bank cases), there's any effect on stakes on knowledge attributions (in the psychological sense which is measured through these simple surveys). So there's no stakes effect for normative facts to mediate.

Now I said above that the inference from stakes to normative facts is one that can be endorsed through philosophical reflection. It does not correspond to an inference that ordinary agents will spontaneously make in the sense that we can expect agents to appeal to it in responding to traditional bank cases. In this sense, stakes don't play a psychological role in ordinary knowledge attributions.

There's another, perhaps less standard sense of 'psychological', however, where it can be said that stakes do play a psychological role in ordinary knowledge attributions. This is the psychological process at play when philosophers or experts make the inference from stakes to normative facts. This "higher" psychological process can be combined with the "lower" process which underwrites inferences from normative facts to knowledge attributions resulting in a mixed process which allows one to deduce facts about knowledge from facts about stakes.

9. Pragmatic Factors and Normative Facts

Our study suggests that knowledge attributions are sensitive to normative facts in the sense that it is awareness of normative facts which causally explains asymmetric knowledge attributions for ordinary agents. We now ask the question whether normative facts are themselves pragmatic facts. I think they are, in the relevant sense. That is, I will argue for

PRAGMATIC NORM: The explicitly mentioned normative fact in the modified bank cases (whether or not it's ok for Bob to assume that the bank will be open on Saturday) is a pragmatic factor (as opposed to a traditional epistemic factor).

Our study suggests, therefore, that ordinary knowledge attributions are sensitive to pragmatic factors—factors which do not include traditional epistemic notions like truth, justification and belief. If this is true, then any support for non-traditional epistemic theories which might have been lost is regained.

Recall that by a 'normative' fact I mean a fact like 'It is acceptable for x to use p in her reasoning'. To be clear, what makes a fact like this true is a cornucopia of grounding facts which may run the gamut from traditional epistemic notions to physical facts to pragmatic notions. We are not interested in whether these grounding facts are themselves pragmatic. We are interested in the question of whether in the modified bank cases, the *difference* in normative facts across the conditions is only a difference in pragmatic factors. The answer to this question is 'yes, the difference is just one in pragmatic factors'.

To see this, note how the normative facts are made explicit in the modified bank cases. In Modified Low Stakes, it is specified that ‘The stakes are low for Bob and his wife. So based on Bob's evidence, it’s appropriate for them to assume in their reasoning that the bank will be open on Saturday.’ And in Modified High Stakes ‘The stakes are high for Bob and his wife. So based on Bob's evidence, it’s not appropriate for them to assume in their reasoning that the bank will be open on Saturday’. Notice that the appropriateness of making the assumption (the normative fact) is explained by what is at stake, a pragmatic factor. We are not told, for example, that the appropriateness is explained by the evidence or any other traditional epistemic factor. All of this strongly suggests that the pragmatic factor at issue is something that is pragmatic.

10. Conclusion

We began noting that a common assumption made in philosophy was that ordinary attributions of knowledge are sensitive to stakes. This assumption has been thought to support certain theories in epistemology (contextualism and pragmatic encroachment). Recent work in experimental philosophy challenged this initial assumption. I argued that the common assumption needs a more subtle interpretation: ordinary attributions are sensitive to stakes in a metaphysical sense and sensitive to normative facts in a psychological sense. On this more subtle interpretation of the data, support for the non-traditional theories may be regained, at least insofar as those theories do a good job explaining how knowledge may be affected by pragmatic factors including what may be acceptable to assume in one’s reasoning.⁶

⁶ I’d like to thank Matt McGrath and Shyam Nair for helpful comments.

WORKS CITED

- Buckwalter, Wesley (2010). Knowledge Isn't Closed on Saturday: A Study in Ordinary Language. *Review of Philosophy and Psychology* 1 (3):395-406.
- Buckwalter, Wesley & Schaffer, Jonathan (2015). Knowledge, Stakes, and Mistakes. *Noûs* 49 (2):201–234.
- Cohen, Stewart (1999). Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives* 13:57-89.
- DeRose, Keith (1995). Solving the skeptical problem. *Philosophical Review* 104 (1):1-52.
- DeRose, Keith (2005). The ordinary language basis for contextualism, and the new invariantism. *Philosophical Quarterly* 55 (219):172–198.
- DeRose, Keith (2011). Contextualism, contrastivism, and X-Phi surveys. *Philosophical Studies* 156 (1):81-110.
- Fantl, Jeremy & McGrath, Matthew (2009a). *Knowledge in an Uncertain World*. Oxford, England: Oxford University Press.
- Fantl, Jeremy & McGrath, Matthew (2009b). Critical study of John Hawthorne's knowledge and lotteries and Jason Stanley's knowledge and practical interests. *Noûs* 43 (1):178-192.
- Feltz, Adam & Zarpentine, Chris (2010). Do You Know More When It Matters Less? *Philosophical Psychology* 23 (5):683–706.
- Francis, Kathryn ; Beaman, Philip & Hansen, Nat (2019). Stakes, Scales, and Skepticism. *Ergo: An Open Access Journal of Philosophy* 6:427–487.
- Hawthorne, John & Stanley, Jason (2008). Knowledge and Action. *Journal of Philosophy* 105 (10):571-590.
- Jackson, Alexander (forthcoming). Are knowledge ascriptions sensitive to social context? *Synthese* 199 (3-4):8579-8610.
- Knobe, J. (2003). Intentional action and side effects in ordinary language. *Analysis* 63 (3):190-194.

May, Joshua ; Sinnott-Armstrong, Walter ; Hull, Jay G. & Zimmerman, Aaron (2010). Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study. *Review of Philosophy and Psychology*_ 1 (2):265–273.

Nagel, Jennifer (2008). Knowledge ascriptions and the psychological consequences of changing stakes. *Australasian Journal of Philosophy*_ 86 (2):279-294.

Pinillos, Ángel (2012). Knowledge, Experiments, and Practical Interests. In Jessica Brown & Mikkel Gerken (eds.), *Knowledge Ascriptions*_. Oxford University Press. pp. 192.

Pinillos, Ángel & Simpson, Shaun (2014). Experimental Evidence in Support of Anti-Intellectualism About Knowledge. In James Beebe (ed.), *Advances in Experimental Epistemology*_. Bloomsbury Academic. pp. 9-44.

Pinillos, N. (2019). Knowledge and the permissibility of action. *Synthese*_ 196 (5):2021-2043.

Rose, David ; Machery, Edouard ; Stich, Stephen ; Alai, Mario ; Angelucci, Adriano ; Berniūnas, Renatas ; Buchtel, Emma E. ; Chatterjee, Amita ; Cheon, Hyundeuk ; Cho, In-Rae ; Cohnitz, Daniel ; Cova, Florian ; Dranseika, Vilius ; Lagos, Ángeles Eraña ; Ghadakpour, Laleh ; Grinberg, Maurice ; Hannikainen, Ivar ; Hashimoto, Takaaki ; Horowitz, Amir ; Hristova, Evgeniya ; Jraissati, Yasmina ; Kadreva, Veselina ; Karasawa, Kaori ; Kim, Hackjin ; Kim, Yeonjeong ; Lee, Minwoo ; Mauro, Carlos ; Mizumoto, Masaharu ; Moruzzi, Sebastiano ; Olivola, Christopher Y. ; Ornelas, Jorge ; Osimani, Barbara ; Romero, Carlos ; Rosas Lopez, Alejandro ; Sangoi, Massimo ; Sereni, Andrea ; Songhorian, Sarah ; Sousa, Paulo ; Struchiner, Noel ; Tripodi, Vera ; Usui, Naoki ; Vázquez del Mercado, Alejandro ; Volpe, Giorgio ; Vosgerichian, Hrag Abraham ; Zhang, Xueyi & Zhu, Jing (2019). Nothing at Stake in Knowledge. *Noûs*_ 53 (1):224-247.

Schaffer, Jonathan & Knobe, Joshua (2012). Contrastive Knowledge Surveyed. *Noûs*_ 46 (4):675-708.

Shin, Joseph (2014). Time constraints and pragmatic encroachment on knowledge. *Episteme*_ 11 (2):157-180.

Turri, John ; Buckwalter, Wesley & Rose, David (2016). Actionability Judgments Cause Knowledge Judgments. *Thought: A Journal of Philosophy*_ 5 (3):212-222.

