



*Cognitive Ecology:  
The Role of the Concept of Knowledge in our Social Cognitive Ecology*

*2010 Epistême Conference  
University of Edinburgh  
June 2<sup>nd</sup>-4<sup>th</sup>, 2010*

*Programme*

*Wednesday June 2<sup>nd</sup>*

- 11am Registration desk opens, main foyer of Informatics Forum
- 12.30-1pm Coffee/tea & biscuits
- 1-2.15pm Professor Peter Graham (UC, Riverside)  
'Knowledge and the Function of Assertion'  
*Chair: TBC*

ABSTRACT. In 'Meaning, Communication, and Knowledge' John McDowell argued that the function of assertive communication was to transmit knowledge. This thesis relates to the nature of linguistic communication and the "norm" of assertion as well as the epistemology of testimony. I shall examine whether an etiological argument can sustain McDowell's thesis, and further explore the consequences of the etiological argument for debates about the norm of assertion and the epistemology of testimony.

2.15-3pm Split graduate session

| <i>Room One<br/>Chair: TBC</i>  | <i>Room Two<br/>Chair: TBC</i>   |
|---|--|
| <p>Boaz Miller (Toronto)<br/>'Epistemic Communalism Defended'</p> <p>ABSTRACT. Mainstream analytic epistemology regards individuals, rather than communities, as the bearers of knowledge or justified beliefs. A few theories break away from this tradition and offer communal analyses of knowledge (Kusch 2003; Longino 2002; Welbourne 2001). Traditional analytic epistemologists may find such accounts objectionable as they deny traditional assumptions about knowledge, such as that knowledge is a species of belief. An argument for epistemic communalism that makes no such objectionable assumptions is provided by Hardwig (1985; 1988; 1991; 1994). Hardwig claims that the justification of the vast majority of a normal adult individual's beliefs is possessed by the epistemic community to which she belongs.</p> | <p>Georgi Gardiner (Edinburgh)<br/>'Thought Experiments, Practical Explications and the Methodology of Epistemology'</p> <p>ABSTRACT. Teleological approaches to elucidating a concept proceed by asking questions such as 'what is the <i>purpose</i> of the concept?' 'What <i>role</i> has it played in the past?', or 'If we imagine a society without the concept, why would they feel the <i>need</i> to invent it?'. The idea behind the teleological approach is that by examining the function of the concept we illuminate the contours of the concept itself. A relatively new development in epistemology, the teleological approach is intertwined with the growth of social epistemology, as the functions posited for epistemic concepts highlight the social role of epistemic language and thought.<br/>The teleological approach can be contrasted</p> |

Hardwig's argument is the following:

1. Individuals do not and in principle cannot possess evidence for most of their beliefs. Rather, they must trust the testimony of others, and this holds both for lay-expert relationships and expert-expert relationship.
2. Either an individual can know "vicariously"—i.e. without personally possessing evidence for what she knows or without understanding what she knows, or only a community can collectively know.
3. An individual cannot know "vicariously".
4. Therefore, only the community can know.

Several social epistemologists have tried to resist this argument. I provide a novel defence of epistemic communalism by arguing that their objections fail both for separate reasons and one common reason. I support my argument with several examples from scientific practice.

I first distinguish the debate about epistemic individualism and communalism from the debate about reductionism and non-reductionism in testimony. Contemporary reductionists (e.g. Fricker 1994; 1995) and non-reductionists (e.g. Coady 1992) agree that when a hearer possesses defeaters, she does not have a right to believe the speaker's testimony. The bone of contention between them is whether a hearer has a presumptive right to believe a speaker in the absence of positive or negative reasons to do so (Lackey 2008, 168-175). I argue that Hardwig's account is consistent with both views. The debate between communalists and individualists is not about a presumptive right to trust, but about justification. Communalists deny that a hearer's indicators of a speaker's sincerity and competence, which may give her reasons to trust, can amount to sufficient justification for knowledge. I counter objections that try to show that they can.

In a response to Hardwig, Adler (1994, 266) argues that once an individual sufficiently corroborates a testimony with other testimonial sources, she can be justified in believing it. I argue, however, that multiple testimonial confirmations are more of the same. They are as dependent on trust as the original testimony, thus they cannot confer justification on the hearer just as the original testimony could not. Moreover, we encounter many multiply confirmed yet false testimonies, such as urban tales (Webb 2004), and there are empirical reasons to doubt that multiple confirmations converge on the truth in the long run as Adler assumes (Lewenstein 1995).

Attempting to resist communalism, Fricker's (2002) argues that individual scientists individuals have indirect social evidence of their peers competence and commitment to nor of truth-telling, which is sufficient to justify their beliefs gained by their peers' testimony even when they cannot verify it by themselves. Fricker assumes, however, that a scientist's appearing to be sincere and competent is correlated with him actually being so, but it is exactly a scientist's difficulty with personally obtaining sufficient justification for believing that the two are indeed correlated that underpins Hardwig's account. Thus, I argue, her objection leads to

with the 'thought experiments' method, which attempts to elucidate a concept by using various cases to mine our intuitions and to test the intuitive extension against a proposed explicit intension. Various reasons have been proffered for abandoning or downgrading the traditional 'thought experiments' approach in epistemology in favour of a teleological approach. In this paper I argue that each reason given for favouring a teleological approach over a thought-experiment approach is misguided, as either the proponent of the new method hasn't actually identified a weakness in the traditional method, or else the proposed new method suffers from the fault as much as the original. I will then suggest two roles in our epistemic theorising unique to the teleological approach, roles which by its very nature the thought experiments method can't play, and so find a place in epistemology for teleological theorising, and hence for ruminations about the specifically social role of epistemic concepts.

I first survey various teleological approaches and, building on the foundations supplied by Kusch (2009) and Kappel (*forthcoming*), I construct a taxonomy of these new methods. Within the teleological (function-first) approach I distinguish between practical explications (synchronic and factual), historical genealogies (diachronic and factual), and hypothetical genealogies (diachronic and stipulative). I sketch the key features of each variation of the teleological approach, and illustrate with an example of each (Kappel's, Kusch's and Craig's (1990) respectively, all of which emphasise the role of our social cognitive economy, and the social character of our epistemic enterprises). I outline what each suggests about the nature of knowledge. I will not attempt to adjudicate between the various permutations of the teleological method, or indeed raise objections to the method, although hopefully by sketching this taxonomy I clarify the terrain for this debate to occur. After explaining teleological approaches, and outlining examples, I explain the key reasons proffered for adopting the teleological approach in place of the thought experiments approach. Drawing on details of the three examples given during my taxonomy, I explain why they are not well-grounded reasons. The reasons I explore are:

- that the thought experiments method hasn't yielded success. This is the idea that the post-Gettier literature is an open grave of different accounts of knowledge, with their various counterexamples, and no real progress is being made;
- since the thought experiments method doesn't ask or answer questions about *why* the concept has the structure it does, it therefore misses a crucial aspect of epistemological theorising, one which the teleological method confronts;
- the teleological approach yields a better understanding of the concept because, unlike the thought experiments method, it reveals the fundamental structure of that concept, rather than just which conditions are and aren't necessary;
- concerns about the justificatory status of conclusions reached via the thought

scepticism.

Goldman (1999, 171; 2001, 106-7) argues that while laypersons cannot usually evaluate the evidence on which experts base their claims, they can obtain justification for their expert-based belief by collecting external reliability indicators. E.g., while a layperson cannot evaluate the reasoning behind an astronomer's prediction of a lunar eclipse or a meteorologist's prediction of rain, she can see for herself if they come true. While she cannot evaluate the correctness of a physician's diagnosis, she can know if the treatment was successful. Such external reliability indicators constitute individuals' justification of their expert-obtained beliefs. I argue that Goldman overplays the availability and salience of such reliability indicators. Laypersons often do not have such reliability indicators, and an expert's success in treating a patient or fixing a car is not as good as an indicator of truth as Goldman argues.

Last, I argue that even if we grant the last three objections, individuals *qua* individuals cannot possess sufficient multiple confirmations, indirect social evidence, or external reliability indicators to justify most of their testimonial beliefs. E.g., an individual who wants to verify a meteorologist's predictions can only personally observe the weather in her immediate environment. This is anecdotal evidence insufficient for determining the meteorologist's reliability. Only a community of geographically disperse weather observers, whose members must trust each other's testimonies, can collectively have sufficient data to corroborate the meteorologist's predictions. Actual examples from science, such as the difficulty in knowing the efficacy of antidepressants (Brown 2008) or the error rates of prenatal screenings (Seavilleklein 2009), demonstrate this claim. Put differently, the last three objections fail because just as individuals *qua* individuals cannot possess first order justification for most of their beliefs, they cannot possess second order justification for their beliefs as well. Only an epistemic community can collectively possess such justification.

The debate between communalists and individualists has framed justification in evidentialist terms. This may give the impression that this is an in-house debate between internalists that need not concern externalists. I argue that this is not so. First, I observe that externalists should be less motivated to resist communalism, as they define knowledge by reference to factors external to the knowing subject. I further draw a comparison between the way laypersons gain knowledge from experts and Goldberg's (2007, 209-26) externalist account of how cognitively immature children gain knowledge from their parents' testimonies. Based on this comparison, I argue that the justification of most of an individual's testimonial beliefs is jointly located in the epistemic community to which she belongs, and she has knowledge only inasmuch its expert members are *de facto* reliable.

experiments method motivate a move to the teleos-based method.

Once I have explained why those reasons are not good motivations for shifting from the thought experiments method to the teleological method, I then suggest two alternative reasons to embrace the teleological approach. First, looking at the role that a concept fulfils will illuminate how precise we can expect the contours of the concept to be, something that we cannot expect the thought experiments approach to tell us. This will tell us at what level of precision we ought to cease expecting intuitions to distinguish whether a particular case falls under the extension of the concept (for example, whether Barney knows or not in various permutations of Barn Façade County). To explain this point I draw on practical explications of concepts from outside epistemology, and show how they reveal how precise or vague the concept is. Second, I argue that the role of teleological approaches is to increase the variety of data we must take account of when seeking reflective equilibrium when offering an analysis of a concept. I hope that these two positive conclusions go some way towards creating peace between teleological-based and thought experiments-based approaches.

This paper also aims to shed light on debates surrounding the two methodologies, by disambiguating the *project* (whether epistemic concepts can be analyzed into constituent parts or merely be better understood), from the *method* (the traditional extension-first approach or the teleos-based approach). I argue that whilst many writers have thus far conflated the two (by pairing the abandonment of the traditional project of analysis with the need to adopt the teleological approach), I argue that either method can be applied to either project. By distinguishing thus I hope to clear up some confusions in the literature (exemplified by arguments to the effect that a teleological approach is preferable as it is still useable even if knowledge is unanalyzable, a claim obviously redundant in the debate between the two methodologies, once we make the distinction and so see that the thought experiments approach is equally consistent).

3-3.45pm

Split graduate session

| <i>Room One</i><br><i>Chair:</i>  | <i>Room Two</i><br><i>Chair:</i>   |
|---|--|
| <p>Daniel Cavedon-Taylor (Birkbeck)<br/>'Depiction as a Source of Knowledge'</p> <p>ABSTRACT. If all pictures were wiped off the face of the earth it would not just be a tragic loss for us from an aesthetic perspective, but from an epistemic one as well. For pictures are frequently sources knowledge about the things they depict: By seeing the court-room sketch we can know that the defendant has a scar above his left eye, by seeing the map drawn by our friend we can know that it is two left turns to get from the station to the pub, and by seeing Muybridge's photographs of animal movement we can know that horses may gallop with all four of their hooves off the ground. We encounter a myriad of pictures in our everyday lives—on billboards, in newspapers on the internet—and we frequently acquire knowledge from them about what they depict (in what follows I restrict myself to discussing propositional knowledge).</p> <p>It is a mundane fact about pictures that they are sources of knowledge, but the epistemology of pictures is a sorely neglected topic. In this paper I take some steps towards remedying that by investigating what kind of knowledge pictures furnish: testimonial, inferential or perceptual?</p> <p>Pictures are commonly held to divide into two types: photographic ones (including film, video and cinema) and hand-made ones (including paintings, sketches and drawings). What is thought to mark the difference between them is the extent to which each is affected by their maker's beliefs and other intentional states. With a painting, every detail is under the picture-maker's control; paintings are through and through affected by their maker's beliefs about what it is they are painting and its properties. With a photograph, by contrast, what gets depicted is whatever is appropriately causally related to the camera, independent of what the photographer intends that their picture depicts. As Gregory Currie illustrates (1991: 24):</p> <p style="padding-left: 2em;">If the painter were having an hallucination, and so thinking there was a pink elephant in front of him, his painting would display a pink elephant... With the photograph things are different. Because of the "mechanical" way photographs are made, it does not matter what the photographer thinks the object in front of the lens looks like; once the camera is set up and the film exposed, the camera records the scene actually before it.</p> <p>In a similar vein, consider also that accidental depiction is a possibility for photographic, but not hand-made, pictures. The photographer may not notice a particular object in frame when they take their picture and because of this be surprised when later they find that object depicted in their picture. In painting, by contrast, because every detail is under the painter's control, accidents of this sort are not possible. Anything that a painting depicts</p> | <p>Shane Ryan (Edinburgh)<br/>'Social Epistemology: Facilitating Both Knowledge and Scepticism with a Special Focus on the Role of Intuition'</p> <p>ABSTRACT. The affect of a believed knowledge claim is such that there is no purpose in further investigation of the matter directly related to the knowledge claim. The achievement of knowledge is an invitation to build, if possible, upon that knowledge. In its broad social consequences knowledge claims can offer assurance, trustworthy information and even provide meaning. However, while claims of achieving knowledge can bear many positive fruit, there are negative consequences if even some of those claims to knowledge turn out to be false. Such knowledge can become a contagion for a broad body of knowledge with the potential to act like a computer virus and contaminate many other knowledge claims that have depended on or used that false knowledge. A general loss of confidence in all knowledge claims may result.</p> <p>If a general and blunt scepticism prevails at a social level, it becomes all too easy for people making knowledge dependent decisions to avoid making hard choices. On the other hand a critical, responsive scepticism is healthy for an individual, a society and even for a refined understanding of the concept of knowledge <i>à la</i> Wittgenstein. (This is particularly so now with the proliferation of "marketed knowledge"—the selling of "wisdom" in books that promise to give the reader the secret to happiness or products based upon knowledge claims, such as alternative pills that are supposed to improve health). While much is written on the conditions under which someone can be said to know, less has been written on how someone should critically assess claims to knowledge. Issues raised by the latter include when it is worthwhile making such assessments, the relevance of the speaker—their competencies and interests; the accessibility of the subject matter; and the most fruitful means to go about assessing such claims. In this section I will make use of Craig's <i>Knowledge and The State of Nature</i> with regard to when and how we attribute knowledge to others and Lackey's work on testimonial knowledge in <i>Learning from Words</i>.</p> <p>This scepticism can in some way resemble that of Socratic questioning. The legitimation of such questioning can also lead to a better relationship between different aspects of society, in which knowledge doesn't become an honorific term that contributes to the dominance of some groups over others. This critical, but conditional scepticism can foster critical, conditioned knowledge claims—illuminating the individual factors that a particular knowledge claim depends upon. This more fine-grained approach can in turn restrict some contagion and loss of confidence that can result when what was thought to be knowledge turns out not to be.</p> <p>I will argue that an important and useful task for</p> |

it does so because the painter intended that it be depicted; depiction in painting is a matter of intention (Wollheim 1980). For this reason there is little sense in being surprised to find that one's painting depicts something which, at the time one painted it, one did not intend it to depict. The difference is sometimes put by saying that paintings are 'intentional' depictions whilst photographs are 'natural' or 'causal' ones (Costello & Phillips 2009).

This intentional/causal distinction is usually made simply by way of spelling out how photographic and hand-made pictures differ *qua* depictions. Whilst some epistemic consequences have been drawn from the distinction, its real epistemic significance has so far gone overlooked for clues are offered here as to what source of knowledge pictures might be. What is suggested by the above conception of hand-made and photographic pictures is that the former are instances of a transmissive source of knowledge whilst the latter are instances of a generative one.

Hand-made pictures are transmissive in that they can enlarge the number of knowers of some proposition; crucially, they cannot enlarge the number of propositions known. This follows from the idea of hand-made pictures as being fully controlled by their maker's intentional states and as therefore not allowing for accidental depiction. For, lack of accidental depiction in painting entails that any knowledge about a painting's subject that we acquire from looking at such pictures is knowledge already possessed by the painter. This suggests that hand-made pictures are sources of second-hand, testimonial knowledge. I present slippery slope considerations that further support this idea and also argue that a significant upshot here is that the domain of testimony is broader than it is traditionally construed, extending beyond 'tellings' and 'gestures'. I further argue that along some dimensions hand-made pictures are epistemically superior to verbal/written testimony. For, insofar as pictures are interpretable by non-language users, the conditions that must be satisfied for them to be knowledge-transmitting are less cognitively demanding those of verbal/written reports.

Photographic pictures, by contrast, are a generative source of knowledge in that they can enlarge the number of propositions known. This follows from the idea of photographs as insensitive to their maker's intentional states and as allowing for accidental depiction. Consider: Thanks to photographs from the Hubble telescope we have come to know a greater number of propositions about the universe, and before Muybridge's photographs of animal movement we did not know that horses may gallop with all four hooves off the ground. The fact that photographic pictures are generative in this respect and not wholly intentional makes it problematic to consider knowledge acquired from them as testimonial, and instead suggest such knowledge to be either inferential or perceptual. I argue that we ought to err on the side of perception here rather than inference. For acquiring knowledge from photographic pictures does not have the phenomenology of being mediated by inferences from the state of the picture to the state of the things depicted, it seems more direct in character and this

the epistemologist is to examine modes of thinking and approaches to learning likely to facilitate knowledge acquisition—in relation to accepting and assessing testimonial claims—and at the same time signposting fallacies and “bad” modes of thought that could result in a promiscuous acceptance of claims to knowledge or, on the other hand, knowledge impoverishment stemming from an over zealous scepticism. One particularly useful element of this is to examine the conditions under which everyday intuitions regarding knowledge claims work well but also when they don't. For example someone's assessment of another's testimony unduly influenced by *ad hominem* factors or an unwarranted weighting of earlier formed beliefs via the primacy effect. Conceiving of useful rules of thumb that people could employ when they're assessing knowledge claims is an aspect of this work.

|   |  |
|---|--|
| <p>suggests such knowledge is perceptual in nature.</p> <p>More remains to be said about the epistemology of pictures, but if what I say is right we have good reason to think that hand-made pictures furnish testimonial knowledge whilst photographic ones furnish perceptual knowledge.</p> |  |
|---|--|

3.45-4.15pm                      Coffee/tea & biscuits

4.15-5.15pm                      Split session

| <i>Room One</i><br><i>Chair: X</i>   | <i>Room Two</i><br><i>Chair: X</i>  |
|--|---|
| <p>Professor Clayton Littlejohn (SMU)<br/>'Knowledge Means Never Having to Say You're Sorry'</p> <p>ABSTRACT. When you have knowledge and when you have adequate justification, you have every right to believe. Knowledge is more than you need. Justification is just enough. If you have the right to believe, you should have the right to reason from that belief or from the proposition believed, provided you know what follows. While this seems like it must be true for theoretical deliberation, what about practical deliberation? Consider two principles:</p> <p>KA: If <math>S</math> knows <math>p</math>, it is proper for <math>S</math> to treat <math>p</math> as a reason for <math>\Phi</math>-ing provided that <math>S</math> knows that the choice to <math>\Phi</math> is a <math>p</math>-dependent choice.</p> <p>JA: If <math>S</math> justified believes <math>p</math>, it is proper for <math>S</math> to treat <math>p</math> as a reason for <math>\Phi</math>-ing provided that <math>S</math> knows that the choice to <math>\Phi</math> is a <math>p</math>-dependent choice.</p> <p>Suppose that the kind of propriety here is epistemic. It is hard to see how JA could be false. If <math>S</math> oughtn't include <math>p</math> in deliberation that concerns some choice known to be <math>p</math>-dependent, there should be some epistemic norm or other that <math>S</math> violates if <math>S</math> deliberates from <math>p</math> anyway. Since it seems that there's not two stages of epistemic appraisal, one that determines whether <math>p</math> is belief-worthy and a second that determines whether <math>p</math> is basis-for-reasoning worthy, it seems that any proposition justifiably believed is justifiably deliberated from, and so any proposition that ought to be excluded from deliberation on epistemic grounds oughtn't be believed on these same grounds. No counterexamples to JA when it is read as pertaining to epistemic propriety. KA should be free from counterexamples as well as it's a consequence of JA.</p> <p>Suppose that the kind of propriety here is not epistemic, but practical. Here it's not so clear what people want to say. In recent work, some authors (e.g., Gibbons) argue that whosoever justifiably believes <math>p</math> can justifiably (i.e., permissibly) act on <math>p</math>. Suppose I've promised you a gin and tonic and justifiably but mistakenly believe that this glass contains a gin and tonic when in fact it contains petrol. I reason as follows: This is a gin and tonic. It's the last gin and tonic. I've promised you a gin and tonic. I should bring you this gin and tonic. The reasoning that terminates in the belief that I should bring you a gin and tonic consists of premises I have the right to believe and</p> | <p>Professor Philip J. Nickel &amp; S. G. T. Kops (Eindhoven)<br/>'Artificial Testimony'</p> <p>ABSTRACT. Testimony is, roughly speaking, an assertion that can be used as a source of belief. In this paper we consider the epistemic status of artificial testimony, i.e., testimony that has no direct human source. For example, an automated bank machine reports on one's account balance, or a computer-aided diagnostic system tells an oncologist that her patient probably has no cancer cells in his liver. We make two assumptions: first, that it is possible to gain warranted belief and even knowledge from such automated systems, and second, that the psychological mechanisms of acquiring belief from artificial sources are fundamentally similar to those grounding warranted belief formed on the basis of normal human testimony.</p> <p>We discuss two problematic consequences of artificial testimony for theories of testimony. The first problem is about what testimony is, i.e., the definition or concept of testimony. It is often claimed that testimony is by definition something produced intentionally by a speaker. For example, Peter Graham defines testimony as implying the speaker's intention that the hearer believe he has relevant authority, as well as the speaker's belief that his claim is relevant to an unresolved question (1997, 227). Jennifer Lackey says that testimony is produced by "an act of communication ... [by which the speaker] intend[s] to express communicable content" (2008, 28, italics in original). Suppose we take the language of these definitions literally, and suppose computers and machines like those in the examples above cannot literally perform acts or have intentions or beliefs. In that case, these definitions appear to relegate artificial testimony to some other category, such as the category of natural signs (Grice 1971) or the category of "everyday chatter" which Lackey carefully distinguishes from testimony (<i>op cit.</i>, 36). Thus it appears that we must either give different definitions of what testimony is, or allow that these computers and machines can literally perform communicative acts and have attitudes of belief and intention.</p> <p>Second, the phenomenon of artificial testimony appears to have serious consequences for our account of the epistemological justification of testimony. One of the goals of an epistemology of testimony is to explain why it</p> |

a conclusion I have a right to believe. The practical deliberation that results in my giving you the gin and tonic (or forming the intention to give you the gin and tonic) involves just the same elements, so it seems that I have the same right to give you the petrol as I have to believe that I should. Some authors (e.g., Fantl & McGrath, Feldman, Gerken, and Neta) seem to want to refrain from saying this and say only that JA is true when we take the notion of propriety to be epistemic.

The cases that are said to cause trouble for JA on its practical reading don't seem to pose much of a threat to KA on its practical reading. So, if you think that the conclusion of practical reasoning (i.e., the  $\Phi$ -ing) and the conclusion of theoretical reasoning (i.e., the belief that you should  $\Phi$ ) have to have the same normative standing, that's a mark in favor of KA. If, however, you think that JA must be true on its theoretical reading, you either need to deal with the alleged counterexamples to JA on its practical reading or explain how the conclusions of practical and theoretical reasoning could have different normative statuses.

In this paper, I want to do three things. First, I shall argue that practical reason and theoretical reason are unified in such a way that JA could only be true if it is true on both a practical and theoretical reading. Essentially, I shall argue that the reasons that bear on whether to judge and whether to act in accordance with that judgment are going to coincide in such a way that there cannot be differences in the normative standing of practical judgment and the actions they rationalize (i.e., you can be obliged to act against your own judgment, but you cannot be obliged to do this unless you are also obliged to refrain from that judgment in the first place.) Second, I shall argue that given an orthodox account of justification, JA couldn't possibly be true on its practical reading. I side with those writers who think that factual mistakes excuse without obviating the need to justify an action that brings about bad consequences or prevents an agent from doing her duty. It's been said that these cases rest on intuitions that no internalist shares. It is at this point that the internalists often dig in and appeal to an anti-luck intuition (i.e., if you didn't know that, say, your actions would bring about bad states of affairs, that's bad luck rather than something that bears on the normative standing of the action and the attitudes that rationalized it). This won't work. Those who opt for this internalist view of justification and defend JA have to give up 'ought' implies 'can', and if you really have no business bashing your opponents with anti-luck intuitions if you reject OIC. Third, it seems that the argument for KA really is an argument for JA. If knowledge means never having to say you're sorry, the same can be said for justification. What we need is an unorthodox account of justification, one that says that there's little that distinguishes justified belief from knowledge except the normatively insignificant stuff that generate an endless supply of Gettier cases.

is sometimes epistemologically legitimate to rely on the assertions of others as a basis for our own beliefs. Thus our philosophical vindication of relevant principles of epistemic justification should be applicable to artificial speech. Some justifications make essential reference to the communicative acts, intentions, mental states, or moral qualities of a source of testimony. For example, it is sometimes said that a speaker's warranted belief in what she asserts forms part of the justification for others to rely on her assertion. Suppose again that existing computers and artefactual systems cannot be said to believe the truth of what they assert. In response to this, several strategies of response are possible. One possible move is to shift reference to such elements from the artefact or system, to the person or persons who designed, programmed, manufactured, and deployed it. A second possible move is to adopt an epistemology of testimony that makes no reference to the attitudinal states or other agential characteristics of a speaker, such as a reliabilist or strictly evidentialist account of testimonial justification.

There are problems with each of these strategies. The first strategy fails to account for the *independent activity* of the system that produces the artificial testimony. There is often no person aware of specific assertions of an automated system. Hence there is no person in a position to take propositional attitudes toward that assertion. The second strategy is problematic because it assimilates testimonial justification to a natural sign, on a par with any other evidential indication of the truth. This option does not do justice to certain unique features of testimonial justification such as its "second-hand-ness" (McMyler 2007). We suggest a third option, along the lines of the account presented by Tyler Burge (1993, 1998), that focuses on the idea of a *source of rationality* conveyed through language.

5.15-6.30pm Professor David Henderson (Nebraska-Lincoln)  
 ‘More Motivated Contextualism’  
*Chair:* TBC

ABSTRACT. I have argued elsewhere that one significantly improved contextualist accounts of knowledge by drawing on some themes regarding knowledge suggested by Edward Craig. If one recognized that a kind of epistemic gatekeeping for variously engaged communities is central to the point and purpose of the concept of *knowledge*, one would be led to intriguing results. Of course, when the community of contextual concern was engaged with matters of practical moment, one would expect to see the contextually variable demands of the sort that have encouraged contextualists like DeRose and Cohen. More surprisingly, perhaps, when the community of contextual concern is one that serves as a source of information for indefinitely many other communities, when it is what I termed a *general purpose source community*, one (as a contextualist) should expect judgments of sorts that have been championed by invariantists like Williamson. There are some limitations to this account motivated contextualist account. Significantly, it does not adequately respect some of the judgments on which fans of pragmatic encroachment and subject sensitive invariantism (like Fantl & McGrath, or Hawthorne) have focused their attention. Here I want to explore the idea that there are multiple points or purposes motivating the concept of knowledge—and that some are not fully accommodated under my earlier talk of “gatekeeping for an epistemic community”. One is a kind of crediting use that is evident when doing the history of science or ideas that has no bearing on who we feature in our contemporary epistemic community.

7.30pm Dinner, John McIntyre Conference Centre

*Thursday June 3<sup>rd</sup>*

9-9.30am Coffee/tea & biscuits

9.30-10.45am Dr. Finn Spicer (Bristol)  
 ‘The Economics of Knowledge’

ABSTRACT. Knowledge has three distinctive features: knowledge can be possessed by individuals; knowledge has a value to the possessor; and knowledge can be transmitted from one individual to another. These three features suggest that economics can provide a framework for understanding the distribution and flow of knowledge. In this paper I begin this line of enquiry, applying the methods of economic and evolutionary game theory to the spread of knowledge, and drawing some early conclusions for epistemologists.

10.45-11.45am Split session

| <i>Room One</i><br><i>Chair:</i>  | <i>Room Two</i><br><i>Chair:</i>  |
|---|---|
| Professor Stephen Grimm (Fordham)<br>‘Which Knowledge is of Most Worth?’<br><br>ABSTRACT. Acquiring knowledge with respect to certain topics clearly seems more worthwhile than acquiring knowledge with respect to other topics—at least, when we consider things from “a practical point of view.” Thus knowing the truth about whether North Korea has nuclear weapons clearly seems more worthwhile, from a practical point of view, than knowing the truth about how many grains of sand are in some random patch of the Sahara. But how do things look when we consider things “from a purely epistemic point of view?” Do certain items of knowledge still seem more | Dr. Axel Gelfert (Singapore)<br>‘Genealogies of the Concept of Knowledge and Their Cognitive-Ecological Dimension’<br><br>ABSTRACT. The present paper analyses and contrasts different genealogical accounts of how the concept of knowledge may have originated from the cognitive needs of communities of mutually dependent finite reasoners. A typical genealogical account of knowledge invites us to imagine a (fictitious) ‘state of nature’ (Craig 1990), in which a ‘minimal’ version of our extant epistemic practices is exhibited, yet which differs in important ways from our own epistemic predicament. For example, interactions between members of the epistemic |

worth having than others?

It would appear so. Take the knowledge about the grains of sand again, and compare it to the sorts of “pure” knowledge that promises to be gained by the researchers at CERN in Switzerland: for instance, knowledge concerning whether corresponding particles of dark matter exist for every known particle of matter. By most lights, finding out whether this thesis about dark matter is true would seem to be of vastly greater epistemic or intellectual importance than knowing certain truths about the grains of sand, even if no practical benefit were to emerge from either item of knowledge. But what is it that accounts for the greater epistemic interest or importance of certain items of knowledge over others—of knowledge concerning the topics addressed at CERN, for instance, as opposed to knowledge about the grains of sand? What is it, in other words, that makes the one item of knowledge of greater *epistemic* interest or importance than the other?

These are the questions that I will address in this paper, and I will argue that the greater epistemic interest or importance of certain items of knowledge derives from the greater epistemic interest or importance of certain kinds of questions. I will then sketch an account of what it is that makes certain questions of greater epistemic interest or importance than others. These sorts of issues are relevant to the “ecology of knowledge,” moreover, because they help to shed light on which questions are most worthy of our attention, and why—questions of great concern when we are trying to decide how best to dedicate our limited intellectual and economic resources.

community associated with the ‘state of nature’ scenario are thought to be minimally social; the degree of ‘division of epistemic labour’ (Williams 1978) is still comparatively low, and epistemic interdependence chiefly takes the form of (testimonial) one-on-one dependence of an *enquirer* (who desires to gain knowledge regarding a particular states of affairs) on an *informant* (who is thought to be in a position to know). While the state of nature, in this respect, differs markedly from the complex divisions of epistemic labour found in contemporary, socially differentiated societies, the genealogical approach also assumes certain important continuities. For example, it is assumed that individuals in the state of nature, while lacking sophisticated social mechanisms of knowledge ascription and division of cognitive labour, are nonetheless equipped with the same ‘on-board’ sources of knowledge and cognitive capacities—such as perception, inference, and memory—as contemporary human reasoners. As I shall argue, this imposes certain constraints on the heuristic value and explanatory power of genealogical accounts.

One of the first examples of a systematic genealogical account of the concept of knowledge is Edward Craig’s ‘practical explication of knowledge’ (1986), later developed more fully in his *Knowledge and the State of Nature* (1990). Earlier, Bernard Williams, in his *Descartes: The Project of Pure Enquiry* (1978) had deployed a similar style of argument, as the foil for his analysis of the (‘non-genealogical’) Cartesian approach to knowledge. Craig’s account, in particular, has been a major influence on certain positions within social epistemology. As Craig sees it, any community of human reasoners is also a community of *agents*, who depend for their well-being and survival on the formation of true (and relevant) beliefs about their environment. This individual need translates into a need for collective strategies of *pooling information*, given that each reasoner individually would likely miss out on much relevant knowledge. An important strategy for pooling information consists in seeking out ‘good informants’—people who are epistemically well-placed and who will pass on relevant information to us. A community of finite reasoners thus will have a strong incentive to develop reliable ways of appraising and labelling ‘good informants’ as such. This, Craig argues, is what drives the introduction of the concept of ‘knowledge’. As communities and interactions between their members become more complex, those ‘indicator properties’ that marked out good informants in the state of nature (such as individual track records of sincerity and competence) become more elusive and difficult to ascertain; however, since the collective need for successful coordination of action persists, a more general term of appraisal—less overtly testimonial in character—is needed to identify good sources of information: hence the abstract concept of ‘knowledge’, which stands at the end of a process of social differentiation and objectivization—a process that has its roots in the cognitive demands faced by individual enquirers seeking out good informants.

Craig’s genealogical explication of the concept of knowledge has since been fruitfully developed and modified by several authors. Thus, Miranda Fricker

(1998) has pointed out that, adopting the genealogical perspective, one should expect the concept of knowledge to have an ineliminable ethical dimension. As Craig himself puts it, our concept of knowledge has its roots in ‘situations in which human beings treat each other as subjects with a common purpose, rather than as objects from which services, in this case true belief, can be extracted’ (1990: 36). This, Fricker argues, gives rise to the possibility of ‘epistemic injustice’. Bernard Williams, in his *Truth and Truthfulness* (2002) has argued that a genealogical approach on the basis of an imaginary ‘state of nature’ must be complemented by a concern for ‘real genealogy’ (2002: 39). Most recently, Martin Kusch (2009) has developed a defence of the value of knowledge on the basis of an extension of Craig’s account, supplemented by work in the sociology of knowledge. As Kusch argues, a genealogical perspective supports the view that knowledge, and the institutions that sustain it—such as the institution of testimony—must be understood not merely in terms of their instrumental usefulness, but as intrinsic values that are central to our social existence. It is this which confers a value to knowledge that goes beyond the instrumental efficacy and usefulness of mere true belief.

The present paper builds on these recent uses of the genealogical approach, but insists on several significant modifications. Most importantly, it is argued, the notion of an epistemic ‘state of nature’ needs to be reconsidered. In the examples of genealogies of knowledge discussed so far, communities in the epistemic state of nature consist of human reasoners that are individually much like us, in that they are equipped with the same (‘on-board’) cognitive capacities, but differ at the collective level, where they lack the sophisticated mechanisms and processes (such as effective and reliable division of epistemic labour) that characterise highly differentiated human societies. While this ‘holding constant’ of individual cognitive capacities has some legitimacy, for example when it comes to explaining the (historically recent) emergence of epistemic standards and conventions (e.g., relating to accuracy and sincerity) in the sciences—an example explicitly discussed by Kusch (2009)—it imposes artificial constraints on the kinds of questions that a genealogical approach can shed light on. Thus, it has recently been argued (Carruthers 2006) that our very capacity to reason—more specifically: our ability to make inferences to the best explanation—may itself have been driven by the evolution of human speech and our capacity to engage in exchanges of testimony. As Peter Carruthers puts it, whenever we decide between different possible explanations, provided ‘the hypotheses in question are expressed in language, the problem of inferring the best explanation reduces to the problem of deciding which of the candidate sentences to believe in the circumstances’ (Carruthers 2006: 364). However, the problem of deciding between candidate sentences that we find ourselves presented with is, of course, simply a version of the problem of which testimony to accept, and when. (On this point, see Gelfert *forthcoming*.) Other authors have argued that certain puzzles in the epistemology of testimony may be easier to resolve, once one takes seriously the suggestion that

|  |  |
|--|--|
|  | <p>testimony—though not <i>guaranteed</i> to be reliable by any innate natural propensities (as Thomas Reid had famously thought)—may nonetheless be considered a <i>natural kind</i> (Michaelian 2008).</p> <p>By developing these lines of arguments and relating them to the genealogical project of explicating our epistemic practices by providing an explanation of their success (and thereby also vindicating, within limits, such practices as our reliance on testimony), the present paper adds an evolutionary-ecological dimension to the hitherto exclusively social-historical genealogies of knowledge.</p> |
|--|--|

11.45-1pm                      Professor Sandy Goldberg (Northwestern)  
‘Knowledge and the Division of Epistemic Labour’  
*Chair:* Professor Alvin Goldman (Rutgers)

ABSTRACT. We depend on others for much of what we know. After enumerating various kinds of dependencies of this sort, I attempt to draw some tentative conclusions about the division of epistemic labour. I am particularly interested in determining how various theories of knowledge fare in acknowledging these sorts of dependencies, and in seeing why the epistemic division of labour might be valuable for a community of epistemic agents.

1-2pm                              Lunch

2-3.15pm                      Professor Lorraine Code (York University, Canada)  
‘An Ecology of Epistemic Authority’  
*Chair:* Dr. Miranda Fricker (Birkbeck)

ABSTRACT. In this paper I will address the power and putative authority of institutional knowledge, taking as my point of entry an observation from Leo Tolstoy’s *Anna Karenina*, which refers to the activities of a commission of inquiry: “All questions were furnished with excellent answers, and answers not open to doubt, since they were not the product of human thought, which is always subject to error, but were the products of institutional activity.” The issue is about an ecology of epistemic authority in social epistemology, as it is operative in the credentialing certain institutions perform and the interplay of credulity and incredulity they initiate. The larger purpose of the analysis, then, is to examine the pretensions of certain forms of institutional knowledge to function as arbiters of truth, and to consider the capacity of such publicly sanctioned knowledge to enact or prevent epistemic injustice especially in places and populations where people are peculiarly vulnerable, epistemically.

3.15-4.15pm                      Split session

| <i>Room One</i><br><i>Chair: X</i>  | <i>Room Two</i><br><i>Chair: X</i>   |
|---|--|
| <p>Professor Mike Ridge (Edinburgh)<br/> ‘How to be an Epistemic Expressivist’</p> <p>ABSTRACT. In this paper, I examine the arguments for moral expressivism to see whether they carry over to the epistemic case. Interestingly, they seem to me to carry over much better at the level of ‘all things considered’ judgments in the epistemic case than at the level of judgments about <i>pro tanto</i> epistemic reasons for belief. I try to follow this argument where it leads, and see whether a hybrid view, according to which epistemic expressivism is true about the ‘all things considered’ but not the <i>pro tanto</i>, is intelligible or even plausible. I then explore the issue of why we might have come to have a</p> | <p>Professor James R. Beebe (Buffalo)<br/> ‘Surprising Connections Between Knowledge and Action: The Epistemic Side-Effect Effect’</p> <p>ABSTRACT. Consider the following, widely endorsed theses: (i) Whether a true belief counts as knowledge depends only upon epistemic factors such as evidence or reliability. (ii) Because the target of philosophical analyses of knowledge is the ordinary person’s concept of knowledge, such analyses should be answerable to data about “what the ordinary person would say” in response to various epistemological thought experiments. In a recent series of experiments my collaborators and I have</p> |

|  |  |
|--|--|
| <p>set of epistemic concepts with this somewhat unexpectedly bifurcated shape.</p> | <p>found that the knowledge attributions of ordinary subjects are influenced in surprising ways by a number of non-epistemic factors. These results call into question the conjunction of (i) and (ii). Our studies reveal that the tendency of ordinary subjects to count a true belief as knowledge is influenced by the following factors: (a) the ethical goodness or badness of an action undertaken in light of the belief, (b) the aesthetic goodness or badness of such an action, (c) the prudential goodness or badness of such an action, (d) the ethical goodness or badness of the attitude the believer takes toward the action s/he performs in light of the belief, and (e) whether an action undertaken in light of the belief violates any salient social norms.</p> <p>Our results reveal unforeseen connections between knowledge and action in folk epistemological practice. Several prominent epistemologists (e.g., Hawthorne, Stanley, Fantl, McGrath) have recently articulated accounts of the relation between knowledge and action. Most of them are committed to (ii). None of them, however, could have predicted that evaluative properties of actions performed in light of beliefs could affect folk assessments of the epistemic status of those beliefs. Furthermore, none of them would allow that these non-epistemic factors <i>should</i> affect the epistemic status of those beliefs. Our experimental results thus pose a challenge to a number of projects in contemporary epistemology.</p> |
|--|--|

4.15-4.45pm                      Coffee/tea ☞ biscuits

4.45-6.15pm                      *Keynote address*  
 Professor Martin Kusch (Vienna)  
 ‘Knowledge and Certainties in the Epistemic State of Nature’  
*Chair:* Professor Duncan Pritchard (Edinburgh)

ABSTRACT. This paper approaches “social cognitive ecology” through the perspectives provided by the “state-of-nature epistemology” or “genealogy of knowledge” proposed by Edward Craig and Bernard Williams. The first half of the paper summarises this proposal and develops it in a number of directions. The second half relates this approach to Ludwig Wittgenstein’s work on “common-sense certainties”. The aim is to show that several of Wittgenstein’s difficult or even obscure claims—e.g., that common-sense certainties cannot be known or doubted—fall into place as truisms when interpreted on the basis of (a revised version of) the genealogy of knowledge.

7pm                                      Wine reception, Playfair Library Hall, Old College

7.30pm                                  Conference dinner, Playfair Library Hall, Old College

*Friday June 4<sup>th</sup>, 2010*

9-9.30am                              Coffee/tea ☞ biscuits

9.30-10.45am                      Professor Alan Millar (Stirling)  
 ‘The Epistemological Significance of Practices’

ABSTRACT. Practices are conceived as essentially rule-governed activities or clusters of such activities. The normative dimension of practices is explained in terms of the idea that participants in practices are subject to, and incur a normative commitment to following, the rules governing those practices. Reasons are given for thinking that there are practices with such a normative dimension. The ensuing discussion explores their epistemological significance. A central theme is that grasp of what participation in a practice requires of participants enables us to understand people within the practice, and make reasonable predictions regarding what they will do, with little if any knowledge of their beliefs and intentions, and behavioural characteristics, beyond what is implied by their being recognizable as participants. The discussion will be based on, but will take further, the discussion of practices in my *Understanding People: Normativity and Rationalizing Explanation* (Oxford UP, 2004).

10.45-11.45am

Split session

|  |   |
|--|---|
| <p><i>Room One</i><br/><i>Chair: X</i></p>   | <p><i>Room Two</i><br/><i>Chair: X</i></p>  |
| <p>Dr. Mikkel Gerken (Copenhagen)<br/>'Epistemic Assessment'</p> <p>ABSTRACT. Knowledge accounts of action and assertion are sometimes said to be motivated by the fact that they align with natural assessments of action and assertion in ordinary language. Competent and rational speakers normally use 'knowledge' and its cognates when they assess action and assertion. In contrast, competing accounts in terms of warrant, evidence or reliability do not straightforwardly align with ordinary language assessments.</p> <p>In response to this line of reasoning, I will set forth a principle that, once it is conjoined with a warrant account of action, can explain the prominence of 'knowledge' in ordinary assessments of action. If this explanation is successful, a central rationale for adopting knowledge accounts of action and assertion is undermined. Moreover, a methodological lesson about the significance of ordinary language for epistemological theorizing may be learned from the explanation.</p> | <p>Dr. Christoph Kelp (Leuven)<br/>'What's the Point of 'Knowledge' Anyway?'</p> <p>ABSTRACT. In <i>Knowledge and the State of Nature</i> Edward Craig defends the thesis that the job of the concept of knowledge is to flag good informants. In recent years, Craig's "good informant thesis" (GIT) has enjoyed an increasing degree of popularity among a number of epistemologists—notably John Greco, Duncan Pritchard, Miranda Fricker and Martin Kusch—who have appealed to it in order to support or motivate a variety of further epistemological claims. Surprisingly, sympathisers of GIT rarely pause to provide independent support for it and tend to refer to Craig's pioneering work instead. Given that this is so, of course, whether or not their wider epistemological goals are achieved will depend heavily on whether or not Craig's case for GIT is itself convincing.</p> <p>In this paper I aim to show, first, that Craig's argument for GIT does not stand up to closer scrutiny and, second, that there is an alternative thesis that retains the virtues of GIT whilst being able to avoid its vices. In order to achieve this I first outline Craig's somewhat unorthodox defence of GIT. From the outset Craig takes GIT to be an initially plausible hypothesis about the job of the concept of knowledge. He then ventures to mount a case in favour of it by arguing, first, that the concept of good informant approximates the concept of knowledge in intuitive intension and extension and, second, that the remaining differences in the intuitive intensions and extensions between the two concepts can be explained away by a process called "objectivisation". According to Craig, then, the concept of knowledge is the objectivised version of the concept of good informant. This is to provide the crucial evidence that completes the case in favour of GIT.</p> <p>My discussion of GIT focuses first on the process of objectivisation and Craig's attempt to put it to use in order to explain away differences in the intuitive intensions and extensions between the concepts of knowledge and good informant. Against Craig, I identify a new type of case in which the intuitive intensions and extensions of the two concepts come apart. Roughly, in this type of case, someone is a good informant because the processes leading him from belief to assertion will reliably lead him to assert only the truth on a certain</p> |

|  |  |
|--|--|
|  | <p>matter whilst, at the same time, intuitively, she isn't a knower because the processes leading to the formation of her beliefs on the matter are highly unreliable. I go on to show that objectivisation as understood by Craig will not serve to explain away the difference in the intuitive intensions and extensions between the two concepts here. The type of case I identify thus provides disconfirmation of GIT.</p> <p>In the second part of my argument against GIT, I consider an alternative thesis concerning the concept of knowledge, the "settled questions thesis" (SQT): the job of the concept of knowledge is to flag settled questions. I show that a version of Craig's argument can be run with SQT in play instead of GIT and that SQT can make sense of the crucial phenomena that provided evidence for GIT. Importantly, however, the application conditions for the concept of a settled question, as opposed to ones for the concept of good informant, are insensitive to the processes leading from an agent's belief to assertion. Since the problem that arose for GIT traded precisely on the fact that the application conditions for the concept of good informant are sensitive to these processes, SQT can avoid this problem. SQT can thus claim the virtues whilst being able to avoid the vices of GIT. Therefore SQT must be regarded as superior to GIT.</p> <p>Importantly, this result also shifts the burden of proof on the shoulders of those who have aimed to employ GIT for wider-ranging epistemological purposes. They must either provide an independent argument in favour of GIT, one that also establishes that GIT is preferable to SQT after all, or show that SQT serves their epistemological purposes just as well as does GIT, or else look for alternative ways to achieve their further epistemological aims.</p> |
|--|--|

11.45-1pm                      Professor Ram Neta (UNC, Chapel Hill)  
    'Knowledge as a Natural Norm of Belief'  
    *Chair*: Professor Klemens Kappel (Copenhagen)

ABSTRACT. According to Ernest Sosa, knowledge involves apt belief, and a belief is apt just in case it is accurate because adroit. But how should we understand the relation between adroitness and accuracy that is here denoted by "because"? I argue that a good answer to this question is suggested by Edward Craig's account of the usefulness of the concept of knowledge.

1-2pm                              Lunch

2-3pm                              Split session

|  |   |
|--|---|
| <p><i>Room One</i><br/><i>Chair</i>: X</p>   | <p><i>Room Two</i><br/><i>Chair</i>: X</p>  |
| <p>Professor Jessica Brown (St Andrews)<br/>         'The Social Role of Knowledge in Practical Reasoning and Intuitive Epistemic Judgements'</p> <p>ABSTRACT. This paper examines the social role of the concept of knowledge in the assessment of action and how this social role affects our intuitive epistemic judgements. It has been recently proposed that</p> | <p>Dr. Gloria Origgi, (CNRS, Institut Nicod)<br/>         'Knowledge, Trust and Epistemic Vigilance'</p> <p>ABSTRACT. Knowledge is a collective good sustained by a thick bundle of trust relations. It has become a platitude in social epistemology to acknowledge the role of trust in creating and sustaining knowledge as a form of distributed cognitive order (cf. Harwdig, 1985; Shapin</p> |

knowledge is the norm of practical reasoning: one is a good enough epistemic position to rely on  $p$  in one's practical reasoning if and only if one knows that  $p$  (e.g., Hawthorne and Stanley). On the knowledge norm view, an ascription/denial of knowledge may be used in the social defence or criticism of our actions. For instance, a mother may criticise her teenage son by saying "You shouldn't have left the party so late. You didn't know that there would be a bus at that time." If knowledge is the norm of practical reasoning, then we might expect ordinary attributions of knowledge to be sensitive to the stakes for the subject of the knowledge attribution. For, it seems plausible that the epistemic position required to act on a certain proposition depends on the stakes: the higher the stakes, the stronger an epistemic position one needs. For instance, I need more evidence to act on the proposition that this curry does not contain peanuts if my child has a potentially life-threatening allergic reaction to peanuts than if she merely prefers not to eat them. Until recently, it was conventional philosophical wisdom that attributions of knowledge are sensitive to the stakes. This was thought to be revealed by our intuitive reactions to contextualist cases, such as DeRose's famous bank case. In the case, DeRose is considering whether to wait in the long Friday queue at the bank to deposit his pay cheque, or go on Saturday. He is stipulated to be in the very same epistemic position in two situations which differ in the stakes. In both, he believes that the bank will be open on Saturday on the basis of the same evidence, namely his previous visit to the bank just two weeks ago. However, in the low stakes scenario, nothing much turns on whether the bank is open on Saturday, whereas in the high-stakes scenario, a huge amount rests on whether the bank is open on Saturday. According to the standard philosophical view, we have the intuition that DeRose knows that the bank is open in the low stakes scenario, but not in the high-stakes scenario. Thus, intuitions about knowledge attributions seem to fit with the role of knowledge in practical reasoning. These two ideas—that knowledge ascriptions are sensitive to the stakes and that knowledge is the norm for practical reasoning—can be used to motivate "impurism", the claim that whether a subject knows that  $p$  depends not only on such traditional factors as whether the subject truly believes that  $p$  and her evidence, but also on the stakes for the subject (see Hawthorne, Stanley, Fantl & McGrath).

Experimental philosophers have recently challenged the claim that folk attributions of knowledge are sensitive to the stakes (e.g., Buckwalter, Feltz & Zarpentine, May, Sinnott-Armstrong, Hull & Zimmerman, and Knobe & Schaffer). They claim to have discovered that folk attributions of knowledge are not sensitive to the stakes and take this to undermine the philosophical position of impurism. One could respond to their argument either by questioning the results and/or their alleged philosophical consequence. Here I focus on the latter response, arguing that the alleged result that knowledge ascriptions are insensitive to the stakes is compatible with leading versions of impurism and the idea that knowledge is the norm for practical reasoning.

Impurism is a metaphysical view that the truth of a knowledge attribution depends not only on

1994; Origg 2004). It is less common, though, to precisely define our trust in others: whether it is based on social and moral values, or on cognitive capacities, or else on social and institutional devices designed to reinforce our reliance in the cognitive order of the society.

In this paper, I will argue that the *stance of trust* (cf. Origg 2007; Holton 1994) we take to gain benefits from other people's knowledge is neither a form of blind trust, nor a default social norm, but a rich, inferential and communicative competence that allows us to sort communicated information. Individuals and institutions that thrive in an information-dense society must display an epistemic vigilant attitude aimed at protecting them from the major risk of accidental or intentional misinformation.

There is moreover no failsafe way to calibrate one's trust in communicated information so as to exactly sort out misinformation. But a number of means that might at least help approximate such sorting exist at the cognitive, social and institutional level. Following Dan Sperber, I call *epistemic vigilance* the suite of mechanisms that allows us to calibrate our trust in others (cf. Sperber *et al.* 2010). Epistemic vigilance, unlike distrust, is not the opposite of trust; it is the opposite of blind trust. Until recently, philosophers and psychologists who have argued that human beings are fundamentally trustful (cf. Burge 1993, Millikan 1984, Gilbert 1993) have claimed that, if people are trustful, they trust transparently, their eyes closed to the possibility of misinformation. In Gilbert's terms, people would be trustful 'by default' and disposed to critically examine communicated information only when circumstances motivate them to do so. This raises the question of how they might recognise such circumstances without being vigilant in the first place.

Mechanisms of epistemic vigilance for filtering information can be found at the level of children's cognitive development (cf. Koenig, Clement, Harris), at the level of comprehension and linguistic interpretation (cf. Graham, 2009) as well at the level of institutional and social filtering of information (cf. Goldman 1999).

Trust is an active stance towards others that makes us epistemic vigilant agents. Accepting the cognitive vulnerability of trusting others implies the cognitive and emotional effort to monitor our informants, evaluate their performance and be sensitive to cues of their trustworthiness. Although our mechanisms of epistemic vigilance may be highly imperfect in sorting the truth, it is their presence at a large scale that entitles us to have a bias for a trustful attitude towards other people's beliefs. To adopt a stance of trust in our interlocutors means to accept a degree of cognitive vulnerability to share with them a number of assumptions, representations and hypothesis that are elicited in communication. We happen to acquire knowledge through communication by actively filtering what we're been told in the context of our own thoughts and epistemic objectives. Appropriateness and credibility vary from one context of communication to another, and may shift within the same conversational setting. But even in the most epistemically relaxed context, there is no transparent belief transfer from a trustworthy speaker to a trustful hearer.

traditional factors but also on the stakes for the subject. This metaphysical claim about the determinants of knowledge does not obviously entail any claim whatsoever about folk intuitions. In general, the truth of metaphysical theories, such as an account of the nature of the big bang, or the mind-brain identity thesis, does not rest on whether the folk have intuitions in line with the true theory. This is so even if the metaphysical theses are stated using terms which the folk understand. For, even if the folk understand the relevant terms, they may be ignorant of the considerations which have led philosophers to endorse the relevant metaphysical thesis.

Despite this, a particular defence of impurism could be hostage to survey work if it uses as a premise the claim that folk attributions of knowledge are sensitive to the stakes. For instance, the following argument for impurism does seem hostage to empirical survey work:

- 1) the folk find it intuitive that, in the low stakes version, DeRose knows that the bank is open on Saturday.
- 2) the folk find it intuitive that, in the high stakes version, DeRose does not know that the bank is open on Saturday.
- 3) the best explanation of 1) and 2) is impurism.

However, impurism would not be so hostage if it is defended by an argument none of whose premises involve the claim that folk intuitions of knowledge are sensitive to the stakes. In the paper, I show how one could motivate impurism by appeal to the idea that knowledge is the norm for practical reasoning without relying on the claim that folk attributions of knowledge are sensitive to the stakes. Of course, this argument is available only if the alleged empirical finding does not undermine the knowledge norm for practical reasoning. I argue that it need not do so. For, one need not adopt a 'curve-fitting' view of the knowledge norm for practical reasoning on which it is an abstraction from folk intuitions about knowledge and practical reasoning. Rather, one may have a more nuanced view on which folk intuitions are some evidence for the norm, without it being an abstraction from them. Alternatively, one may provide a theoretical motivation for the knowledge norm (e.g., Fantl & McGrath)

The argument for impurism in question does not start with claims about folk intuitions about the bank case. Rather, it starts from general principles and argues that it follows from these principles that there are cases like the bank case in which the subject knows that  $p$  in the low stakes scenario, but is not in a good enough epistemic position to rely on  $p$  in her practical reasoning in the high stakes scenario. The argument then combines the existence of such a case with the knowledge norm for practical reasoning to conclude that knowledge depends on the stakes. The argument starts from the claim that a subject may know a proposition yet be able to further strengthen her epistemic position, say by acquiring further evidence. For instance, I may know by sight that a certain bird in the garden is a blackbird, but I could strengthen my epistemic position by checking with an ornithologist. The argument then adds a further principle about practical reasoning, namely that as the stakes rise, one needs a stronger epistemic position to rely on a

Trust is thus an ingredient of knowledge in a social world insofar as mechanisms of epistemic vigilance exist and are actively used by trustful agents in acquiring information.

|   |  |
|---|--|
| <p>proposition in one's practical reasoning. Given these first two principles concerning knowledge and practical reasoning, it follows that we can construct case pairs which differ only in the stakes such that the subject knows that <i>p</i> in the low stakes scenario but is not in a good enough epistemic position to act on <i>p</i> in the high stakes scenario. By appeal to the knowledge norm for practical reasoning, it follows that the subject does not know the relevant proposition in the high stakes scenario. Thus, we have a difference of knowledge across the two scenarios even though they differ only in the stakes for the subject. (This form of argument is found, e.g., in Fantl &amp; McGrath.)</p> |  |
|---|--|

3-4.15pm

Professor Hilary Kornblith (UMass)

'Why Should We Care About the Concept of Knowledge?'

*Chair*: Professor David Bloor (Edinburgh)

ABSTRACT. Why is it important to have a concept of knowledge? It is certainly possible to have knowledge without having the concept of knowledge. Infants do it, and non-human animals do it as well. So what epistemic advantages do we gain by having a concept of knowledge? This paper attempts to answer this question by examining some of the differences between human adult cognition and cognition in infants and non-human animals.

4.15-4.45pm

Coffee/tea &amp; biscuits

4.45-5.45pm

Split session

| <i>Room One</i><br><i>Chair</i> : X  | <i>Room Two</i><br><i>Chair</i> : X   |
|--|---|
| <p>Aidan McGlynn (UT, Austin/Aberdeen)<br/>'Interpretation, Explanation, and Knowledge Maximization'</p> <p>Timothy Williamson has proposed that 'we should replace true belief by knowledge in a principle of charity constitutive of content' (2007: 264; see also his 2000: 267). Understanding charity as knowledge maximization is intended to have an anti-sceptical payload (as is the Davidsonian principle of charity Williamson wishes to replace). On Williamson's account, beliefs of their nature tend to constitute knowledge. So knowledge maximization implies that we tend to apply concepts correctly in judgment, thereby undermining 'judgment scepticism', according to which we systematically misapply everyday concepts such as <i>mountain</i>, <i>person</i>, <i>belief</i>, or <i>knowledge</i> in judgment.</p> <p>Williamson's argument for the principle of knowledge maximization comes in two stages. First, he argues that a crude principle of truth-maximization leads to incorrect assignments of reference, since it ignores the role played by causal connections in fixing the referents of our expressions. Second, he suggests that such causal connections are important in this regard because they act as <i>channels for knowledge</i>. Assignments of reference that respect these causal connections win out because they can maximize knowledge even when they do not maximize true belief. It is vital to distinguish these two stages in Williamson's argument. I do not doubt the</p> | <p>Jason Kawall (Colgate)<br/>'Epistemic Saints, Heroes, and Supererogation'</p> <p>ABSTRACT. In this paper I examine several normative notions that seem to play important roles in ethics—in virtue ethics, in particular—yet which seem to play little role in our epistemic thought. I consider whether there are viable epistemic counterparts to these notions, and whether these counterparts ought to play a similarly important role in our epistemic cognitive ecology (even if they do not, currently). In particular, I focus on saints, heroes, and supererogation—those agents or actions who/that, in some sense, go beyond what is merely required from an epistemic point of view. My primary goal will be to bring out certain difficulties in attempting to extend these notions (supererogation, heroism, saintliness) to epistemic discourse, and to point to some ways of moving investigation of these notions forward. Beyond this, I will begin to sketch how incorporating these notions into our social cognitive ecology might be of particular value in establishing epistemic ideals for individuals, and in characterizing the good epistemic life.</p> <p>Consider, to begin, with how we might attempt to understand epistemic supererogation. There are very few discussions of the notion. Plantinga, in his <i>Warrant: The Current Debate</i>, uses the term in passing while arguing against deontological accounts of epistemic justification. For Plantinga, epistemic supererogation seems to consist in an effort to fulfil one's duties that goes beyond what is</p> |

importance of respecting causal connections when engaged in interpretation, and so Williamson is surely right to hold that a principle of charity that would have us ignore such connections is problematic. But this is an insight that one can appreciate without making the further step to knowledge maximization.

My argument for this latter claim proceeds in three steps. First, I note that given the way Williamson defends knowledge maximization against certain objections, it starts to seem as if what is of fundamental importance in Williamson's examples is not the maximization of knowledge, but rather that the speaker be causally connected to the right objects 'in the right way for reference, whatever that is' (2007: 263). Second, I argue that, contrary to what Williamson suggests (2007: 264), there are cases in which a causal connection is *not* a channel to knowledge. For instance, one may learn that there is a ninety percent chance one has swallowed a powerful hallucinogenic drug. In fact, one hasn't taken the drug. One sees an object, and forms judgments about its shape and colour on that basis. It is plausible that one's reasons for thinking it very likely that one has swallowed the drug act as a defeater, preventing any of the judgments about the object one makes from constituting knowledge. The object is the most eligible candidate to be the referent of any pronouns one uses in expressing the judgments one has formed, but in this case it is clear that this is not because assignments that make it the referent maximize one's knowledge, for they do not. Third, borrowing an idea that is at least implicit in Davidson's own writings, I suggest that the real significance of causal connections lies in the fact that an assignment of reference that respects them enables an interpreter to best *explain* the speaker's linguistic and non-linguistic behaviour. As David Wiggins puts the point (1980: 199), '[i]t is necessary to replace the idea of Charity with an idea that founds interpretation not in maximization but in *explanation* (generously conceived)'. In the final section, I discuss the import this shift has for our engagement with judgment scepticism. Williamson argues that the notion of a 'better explanation' may not be available to us if we wish to avoid simply begging the question against the judgment sceptic (2007: 224). If this is correct, the upshot may be that debates about the theory of interpretation lack the significance for our evaluation of scepticism that Davidson and Williamson wish to attribute them.

epistemically required—to be exceptionally vigilant in seeking the truth. Plantinga's basic point is that even if an agent were to behave in such an epistemically supererogatory fashion, this would not yet guarantee that such an agent would acquire knowledge. But more broadly, a number of possible understandings of epistemic supererogation soon come to mind:

- (1) e.s. (epistemic supererogation) as acquiring more justification or warrant for a given belief than is needed in order to acquire knowledge in a given situation. (Perhaps checking more sources than is necessary, running an exceptional number of tests to be sure that a result is correct, etc.)
- (2) e.s. as attempting to acquire more knowledge than is epistemically expected or required. The student who studies and tries to learn far more than is required or expected in a given course.
- (3) e.s. as involving some sacrifice of epistemic goods for the epistemic benefit of others, beyond what is required. (E.g., a professor devoting more time to teaching an intro-level course than is required of a competent teacher in order to ensure that students gain more knowledge in a subject area, while sacrificing her own research. Compare: devoting one's time to helping others beyond what is morally required, even at the expense of one's own well-being).
- (4) e.s. as involving some sacrifice of goods in general (not only epistemic goods) for the sake of achieving a great epistemic good—for herself or for others. (The scientist who loses out on a social event in order to work in her lab; or a devoted teacher who sacrifices her afternoon in order to have time to help students, etc.)

Of course there are other viable possibilities. I hope to begin an engagement with these issues, with an eye towards how useful these concepts might be to our social cognitive ecology.

Similarly, we can consider what it might be to be an epistemic saint or hero, and again many possible understandings come to mind. For example, we could draw on J. O. Urmson's classic discussion in his 'Saints and Heroes', and hold that an epistemic saint is one who goes beyond her epistemic duties (if any) in her actions, either effortlessly, or by control of contrary inclinations. Thus we might think again of a scientist who devotes herself to her lab over the course of years, sacrificing other interests. Following Urmson's characterization of the moral hero, we could hold that an epistemic hero is one who goes beyond her epistemic duties (if any) in her actions, through control of fear. We might imagine a researcher working in highly dangerous conditions—perhaps in the midst of a war-zone, or studying storms from within, and so on. More broadly, we might think of epistemic saints as individuals who consistently acquire more justification for their beliefs

|  |   |
|--|---|
|  | <p>than is required for knowledge, and so on.</p> <p>As we consider these neglected notions, and possible characterizations of them, we are soon drawn into questions concerning the nature of the good epistemic life, and it is here that I think these notions can be most helpful, and worth integrating into our epistemic discourse (and broader cognitive ecology). As epistemic agents, we do not simply engage in a series of isolated events of belief-formation. We lead rich lives, with competing values, and limited time. I argue in this paper that we would do well to begin to articulate what a good epistemic life would involve, and to explore more fully our epistemic ideals. Who should we look to as agents who have led exemplary epistemic lives? What are exemplary epistemic acts? Answers to such questions will, I believe, be of great use to us as we shape our epistemic lives, both as individuals, and as societies.</p> |
|--|---|

5.45-6.45pm

Split session

| <i>Room One</i><br><i>Chair: X</i>   | <i>Room Two</i><br><i>Chair: X</i>  |
|--|---|
| <p>Dr. Matthew Chrisman (Edinburgh)<br/>‘The Concept of Knowledge and the Aspectual Classification of ‘Knows’</p> <p>ABSTRACT. It’s commonplace in epistemology to refer to the <i>states</i> of belief and knowledge. However, from the point of view of understanding the normativity or value involved in the concept of knowledge, it’s tempting to treat of belief as a sort of cognitive performance, activity, or accomplishment, which, when all goes well, counts as knowledge. In this paper, I use aspectual data from the semantics of ‘knows’ to confirm the idea that the concepts of belief and knowledge are state-concepts. Then I explore an alternative strategy for accounting for the normativity or value involved in the concept of knowledge. This involves the notion of ‘state norms’, whose validity admits of an ecological form of explanation linking them to the epistemic communities they constitute.</p> | <p>Professor Sarah Wright (Georgia)<br/>‘Knowledge and Social Roles: A Virtue Approach’</p> <p>ABSTRACT. Consider two claims: “My doctor knows that drug x is safe and effective” and “My neighbour knows that drug x is safe and effective.” Why does my learning the former make it more rational for me to take drug x than my learning the latter? If knowledge provides a set standard for all cognizers to meet, this is inexplicable. Cases like these show that our concept of knowledge is sensitive to the social roles of those to whom knowledge is ascribed. When we say that a doctor knows a medical fact we are saying something stronger than when we say that a non-doctor knows that same fact. How can we explain this feature of our concept of knowledge?</p> <p>I argue that the sensitivity of knowledge to social roles is not well explained by standard contextualist approaches. Such approaches focus too strongly on the conversational context in which a knowledge claim is made, and in so doing, leave no room for the more stable social roles of the targets of knowledge attribution. Thus contextualism cannot explain the differences between the two claims above; if both are expressed within the same context, the same contextually-determined standards apply to both knowledge attributions, and contextualism still makes the difference between them inexplicable.</p> <p>Instead we should turn to virtue epistemology, in which social roles find a natural home. Just as the demands of moral virtue depend in part on the social roles of those who possess those virtues, the demands of intellectual virtues also depend on the social roles of those who possess them. Courage in a soldier will require far more than courage in a child preparing to walk past the class bully. However, both can appropriately be said to possess the virtue of courage and to display that courage in their actions. This dependence of virtues on social roles explains why the standards for knowledge also depend on social roles. Virtue epistemologies will</p> |

|  |   |
|--|---|
|  | <p>insist that possession of intellectual virtue (or something quite near to it) is required for knowledge; if the standards for possession of intellectual virtue are higher for someone because of their social role, then the concomitant standards for knowledge are higher as well.</p> <p>But this is only one of two ways that knowledge depends on social roles in a virtue epistemology approach. The second connection is supplied by the credit accounts of knowledge offered by John Greco, Wayne Riggs, and Linda Zagzebski. Each proposes minor variants on the idea that a one knows in a particular instance only if the true belief in question is the result of one's believing virtuously. It is the connection of "being the result of" believing virtuously that makes the true belief creditable to the believer. But how shall we explain the relation of "being the result of"? Greco has forcefully argued that this relation itself depends on the context of attribution. I will show how social roles also play a role in setting this context, and so play a role in attributions of credit.</p> <p>These two relations between knowledge and social roles together provide an explanation for the rationality of being willing to take drug x after learning that "My doctor knows that drug x is safe." My decision procedure concerning my healthcare choices should be appropriately sensitive to the social roles of those to whom knowledge is ascribed. To ignore these social roles would be to irrationally ignore the social character of knowledge ascriptions.</p> |
|--|---|

6.45pm

Conference end

7pm

Dinner, venue TBC